CAUSAL NETWORKS
CAUSAL MODELS

Fabio Stella
Department of Informatics, Systems and Communication,
University of Milan-Bicocca
Viale Sarca 336, 2016 Milan, ITALY
e-mail: fabio.stella@unimib.it
Twitter: FaSt@FabioAStella

In this lecture you will learn about fundamental concepts and notations needed to clearly describe causal models.

In particular, the lecture presents and discusses the following:

- do-operator
- Observational and interventional study/data
- Pre-intervention and post-intervention distribution
- Modularity
- Backdoor criterion and adjustment

PART I

THE DO-OPERATOR AND INTERVENTIONAL DISTRIBUTIONS

Causal models are essential for identification of causal quantities.

From the **IDENTIFICATION-ESTIMATION FLOWCHART**

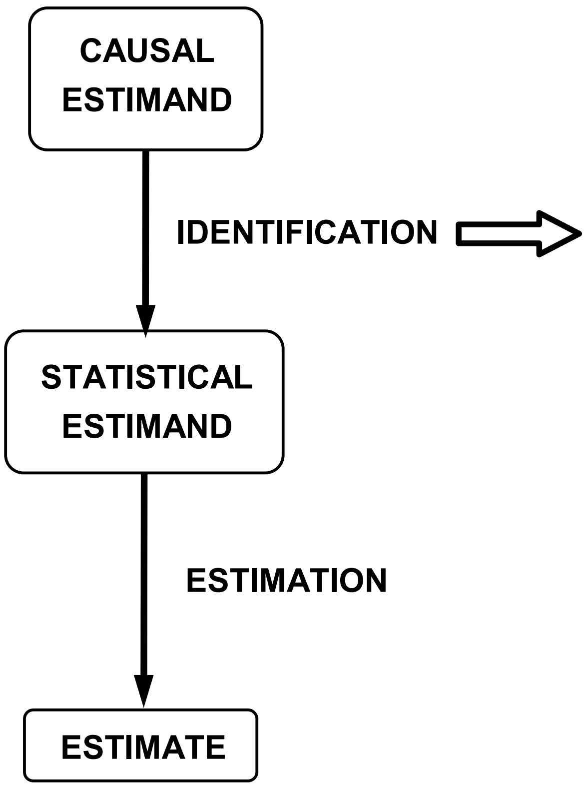

the process of moving from a causal estimand to a statistical estimand

To do this step we must have a causal model.

We update to the following version of the **IDENTIFICATION-ESTIMATION FLOWCHART**

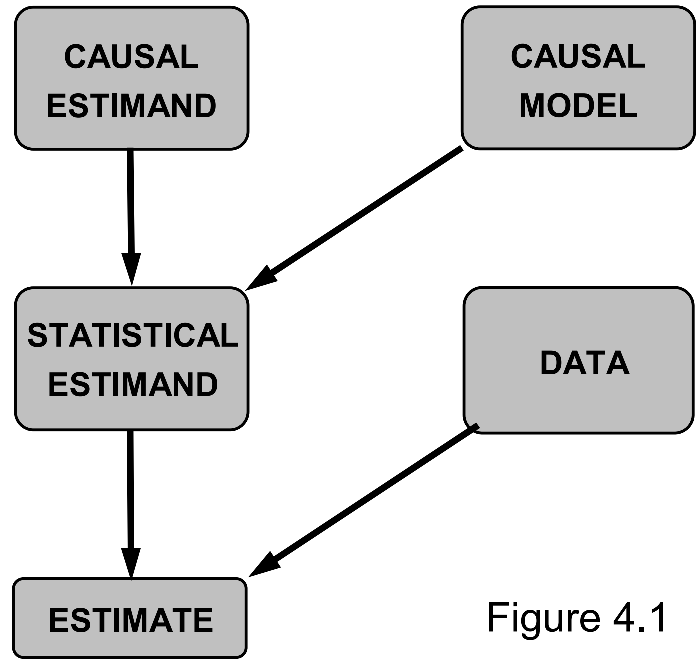

Figure 4.1

In this lecture we explain how to identify causal quantities and formalize causal models.

When we collect data on factors associated with wildfires, we are actually searching for something we can **INTERVENE** upon in order to decrease wildfire frequency.

When we perform a study on a new cancer drug, we are trying to identify how a patient's illness responds when we **INTERVENE** upon it by medicating the patient.

When we research the correlation between violent television and acts of aggression in children, we are trying to determine whether **INTERVENING** to reduce children's access to violent television will reduce their aggressiveness.

THE ULTIMATE AIM OF MANY STATISTICAL STUDIES IS TO PREDICT THE EFFECTS OF INTERVENTIONS.

As you have undoubtedly heard many times in statistics classes,

"CORRELATION IS NOT CAUSATION"

A mere association between two variables does not necessarily mean that one of those variables causes the other.

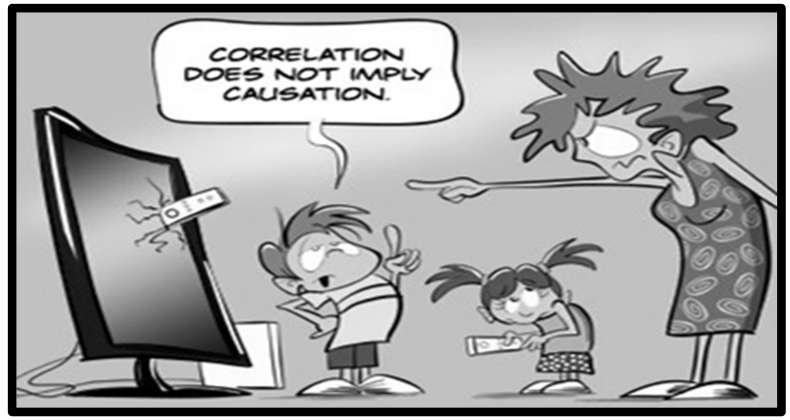

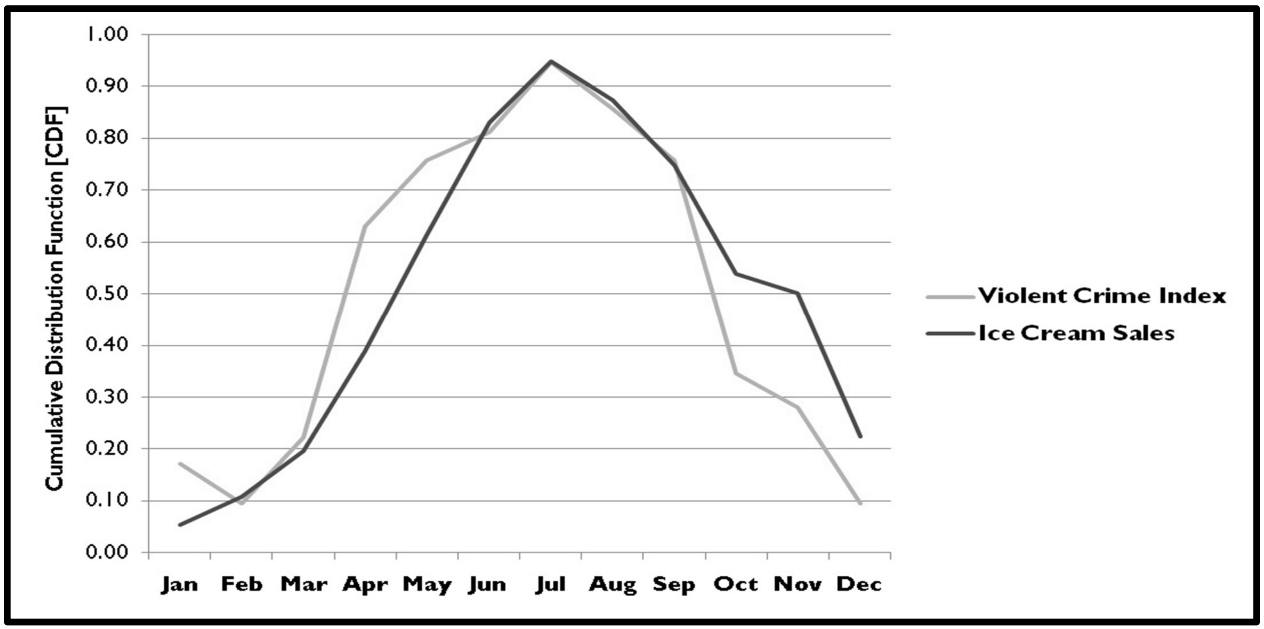

The famous example of this property is that an increase in ice cream sales is correlated with an increase in violent crime—not because ice cream causes crime, but because both ice cream sales and violent crime are more common in hot weather.

For this reason, the **RANDOMIZED CONTROLLED EXPERIMENT** is considered the golden standard of statistics.

In a properly randomized controlled experiment, all **FACTORS** that influence the **OUTCOME** variable are either static, or vary at random, except for one—so any change in the outcome variable must be due to that one input variable (factor).

Unfortunately, many questions do not lend themselves to randomized controlled experiments.

We cannot control the weather, so we can't randomize the variables that affect wildfires.

Even randomized drug trials can run into problems when participants drop out, fail to take their medication, or misreport their usage.

We could conceivably randomize the participants in a study about violent television, but it would be difficult to effectively control how much television each child watches, and nearly impossible to know whether we were controlling them effectively or not.

In cases where randomized controlled experiments are not practical, researchers instead perform **OBSERVATIONAL STUDIES**, in which they merely record data, rather than controlling it (**INTERVENTIONAL STUDIES**).

The problem of such studies is that it is difficult to untangle the causal from the merely correlative.

The difference between **INTERVENING** on a variable and **CONDITIONING** on that variable should, hopefully, be obvious.

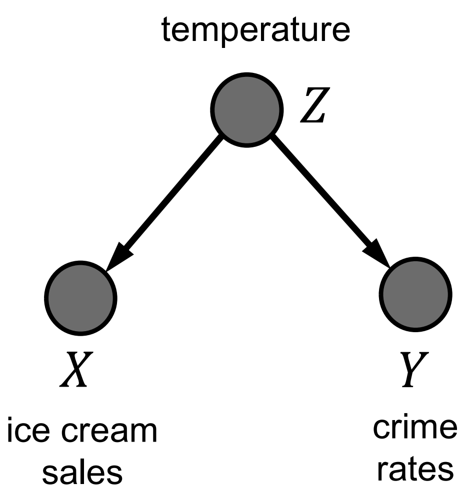

Figure 4.2

Consider, for instance, Figure 4.2 that shows a graphical model of our ice cream sales example, with

- X as ice cream sales
- Y as crime rates
- Z as temperature

INTERVENING on variable X of model M

↓

fix the value of X

↓

we change the system, and the values of other variables often change as a result

CONDITIONING on variable X of model M

↓

we change nothing; we merely narrow our focus to the subset of cases in which the variable takes the value we are interested in

↓

what changes, then, is our perception about the world, not the world itself

In cases where randomized controlled experiments are not practical, researchers instead perform **OBSERVATIONAL STUDIES**, in which they merely record data, rather than controlling it (**INTERVENTIONAL STUDIES**).

The problem of such studies is that it is difficult to untangle the causal from the merely correlative.

The difference between **INTERVENING** on a variable and **CONDITIONING** on that variable should, hopefully, be obvious.

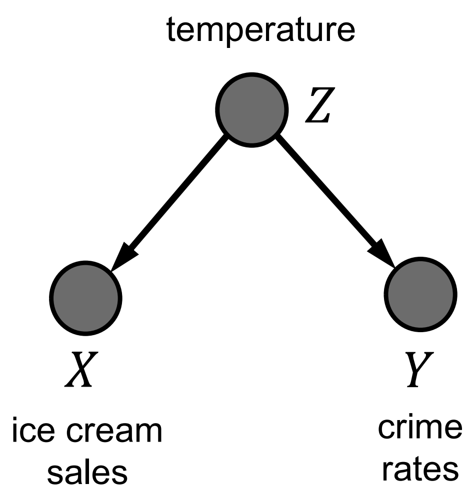

Figure 4.2

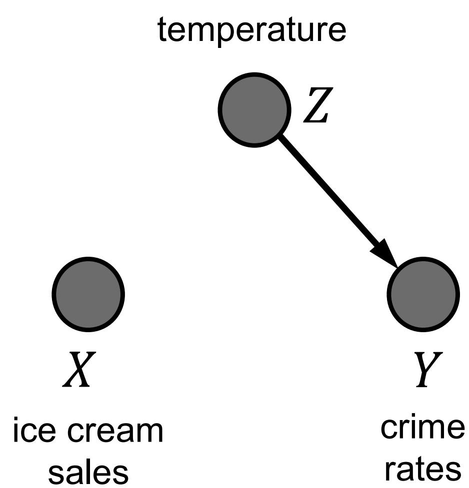

Figure 4.3

When we intervene to fix the value of a variable, we curtail the natural tendency of that variable to vary in response to other variables in nature.

This amounts to performing a kind of **SURGERY ON THE GRAPHICAL MODEL**, removing all edges directed into that variable.

If we were to intervene to make ice cream sales X low (say, by shutting down all ice cream shops), we would have the graphical model shown in Figure 4.3.

When we examine correlations in this new graph (Figure 4.3), we find that crime rates Y are, totally independent of (i.e., uncorrelated with) ice cream sales X since the latter is no longer associated with temperature Z.

In other words, even if we vary the level at which we hold X (ice cream sales) constant, that variation will not be transmitted to variable Y (crime rates).

We introduce the operator to represent **INTERVENTION**.

In the regular notation for probability, we have conditioning, but that isn't the same as intervening.

CONDITIONING on $X = x$ just means that we are restricting our focus to the subset of the population to those who received treatment $X = x$.

In contrast, an **INTERVENTION** would be to take the whole population and give everyone treatment $X = x$.

We denote **INTERVENTION** with the **DO-OPERATOR** $do(X = x)$.

This is the notation commonly used in graphical causal models, and it has equivalents in potential outcomes notation, as follows:

$$P(Y(x) = y) \triangleq P\big(Y = y|do(X = x)\big) \triangleq P\big(y|do(x)\big)$$

The ATE (average treatment effect), when the treatment is binary, can be written as follows:

$$\mathbb{E}[Y|do(X = 1)] - \mathbb{E}[Y|do(X = 0)]$$

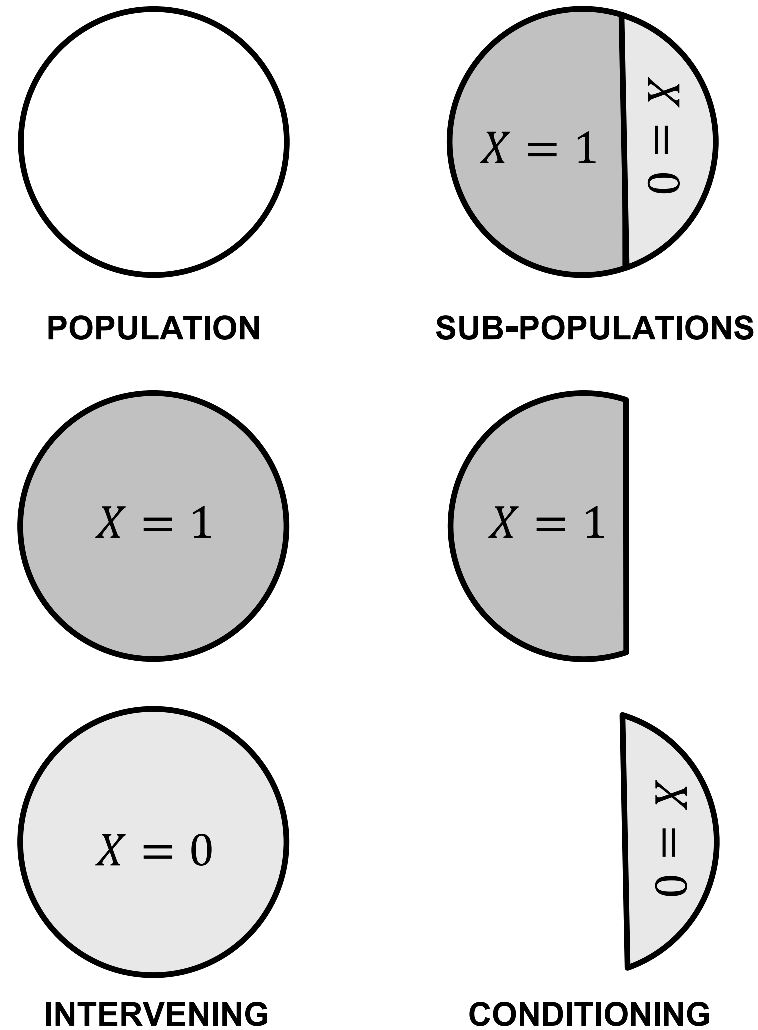

OBSERVATIONAL DATA

OBSERVATIONAL DISTRIBUTION

$$P(Y|X = x) = P(y|x)$$

POPULATION

INTERVENTIONAL DATA

INTERVENTIONAL DISTRIBUTION

$$P\big(Y|do(X = x)\big) \triangleq P\big(y|do(x)\big)$$

$X = 1$

All the units of this subpopulation are treated $X = 1$

$X = 0$

All the units of this subpopulation are not treated $X = 0$

CONDITIONING

All the units of the population are treated $X = 1$

$X = 1$

All the units of the population are not treated $X = 0$

$X = 0$

INTERVENING

An expression Q with do in it is said to be an **INTERVENTIONAL EXPRESSION**.

An expression Q without a do in it is said to be an **OBSERVATIONAL EXPRESSION**.

An interventional expression which can be reduced to an observational expression is said to be **IDENTIFIABLE**.

An **ESTIMAND** is said to be

- **CAUSAL**, whether it does contain the *do*-operator
- **STATISTICAL**, whether it does not contain the *do*-operator

Whenever, $do(x)$ appears in expression Q after the conditioning bar, it means that everything in that expression Q is in the **POST-INTERVENTION WORLD** where intervention $do(x)$ occurs.

$$\mathbb{E}[Y|do(x),\mathbf{Z}=\boldsymbol{z}]$$

Refers to the expected outcome Y in the (**POST-INTERVENTION**) sub-population where $\mathbf{Z}=\boldsymbol{z}$ after the whole sub-population has taken treatment $X=x$.

$$P(Y|do(x),\mathbf{Z}=\boldsymbol{z})$$

POST-INTERVENTION DISTRIBUTION

$$\mathbb{E}[Y|\mathbf{Z}=\boldsymbol{z}]$$

Refers to the expected outcome Y in the (**PRE-INTERVENTION**) population where individuals take whatever treatment X they would normally take).

$$P(Y|x,\mathbf{Z}=\boldsymbol{z})$$

PRE-INTERVENTION DISTRIBUTION

PART II

MODULARITY AND ADJUSTMENT FORMULA

Before we can describe a very important assumption, we must specify **WHAT A CAUSAL MECHANISM IS.**

There are a few different ways to think about causal mechanisms. In the following we let a **CAUSAL MECHANISM** to be a mechanism that generates X_i as the conditional distribution of X_i given its parents (causes) $pa(X_i)$, i.e., the following conditional distribution $P\big(X_i|pa(X_i)\big)$.

The main assumption we need to progress toward **CAUSAL NETWORKS** is that **INTERVENTIONS ARE LOCAL.**

In particular, we assume that intervening on a variable X_i only changes the causal mechanism for X_i; it does not change the causal mechanisms that generate any other variables X_j.

MODULARITY – INDEPENDENCE MECHANISM – INVARIANCE

If we intervene on a set of nodes/variables $\mathbf{S}$, setting them to constants, then for all $X_i \in \{X_1, X_2, \dots, X_n\}$, we have the following:

1. If $X_i \notin \mathbf{S}$, then $P\big(X_i = x|pa(X_i)\big)$ remains unchanged,
2. If $X_i \in \mathbf{S}$, then $P\big(X_i = x|pa(X_i)\big) = 1$, if x is the value that X_i was set to by the intervention $do(X_i = x)$; otherwise $P\big(X_i = x|pa(X_i)\big) = 0$.

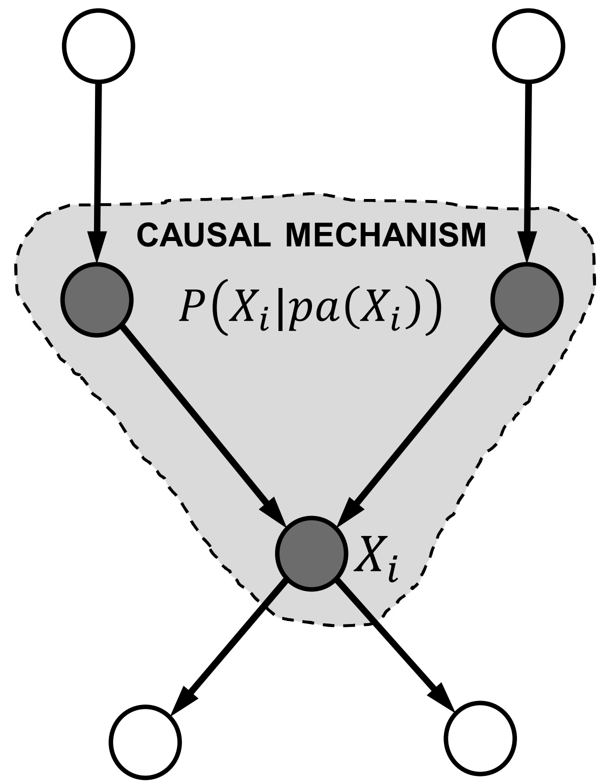

Figure 4.4(a)

We could write condition 2) below as follows:

- $P(X_i = x|pa(X_i)) = 1$ if x is consistent with the intervention
- $P(X_i = x|pa(X_i)) = 0$ otherwise.

In the future we say that, if $X_i \in \mathbf{S}$, a value x of X_i is consistent with the intervention on X_i, if x equals the value that X_i was set to in the intervention, i.e., $do(X_i = x)$.

The causal graph for interventional distributions is simply the same graph that was used for the observational joint distribution, but with all of the edges to the intervened node(s) removed.

MODULARITY – INDEPENDENCE MECHANISM – INVARIANCE

If we intervene on a set of nodes/variables $\mathbf{S}$, setting them to constants, then for all $X_i \in \{X_1, X_2, \dots, X_n\}$, we have the following:

1. If $X_i \notin \mathbf{S}$, then $P(X_i = x|pa(X_i))$ remains unchanged,
2. If $X_i \in \mathbf{S}$, then $P(X_i = x|pa(X_i)) = 1$, if x is the value that X_i was set to by the intervention $do(X_i = x)$; otherwise $P(X_i = x|pa(X_i)) = 0$.

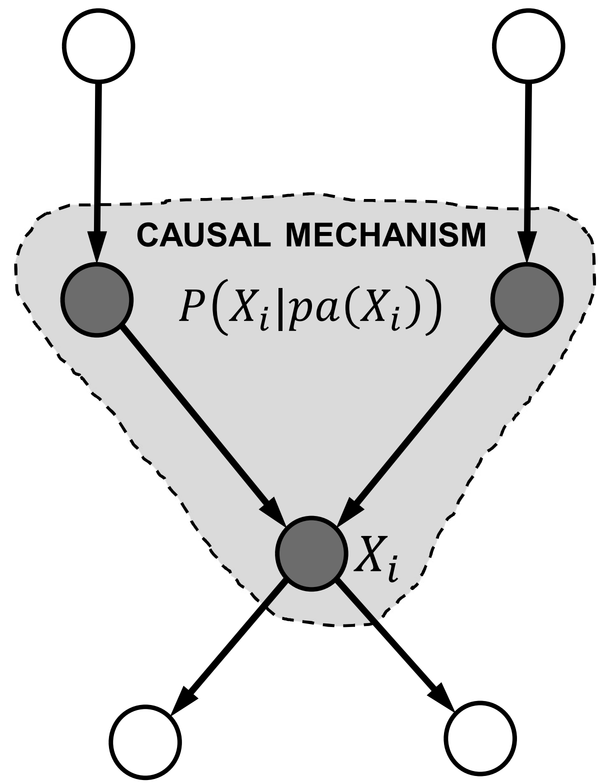

Figure 4.4(a)

We could write condition 2) below as follows:

- $P(X_i = x|pa(X_i)) = 1$ if x is consistent with the intervention
- $P(X_i = x|pa(X_i)) = 0$ otherwise.

In the future we say that, if $X_i \in \mathbf{S}$, a value x of X_i is consistent with the intervention on X_i, if x equals the value that X_i was set to in the intervention, i.e., $do(X_i = x)$.

The causal graph for interventional distributions is simply the same graph that was used for the observational joint distribution, but with all of the edges to the intervened node(s) removed.

MODULARITY – INDEPENDENCE MECHANISM – INVARIANCE

If we intervene on a set of nodes/variables $\mathbf{S}$, setting them to constants, then for all $X_i \in \{X_1, X_2, \dots, X_n\}$, we have the following:

1. If $X_i \notin \mathbf{S}$, then $P(X_i = x|pa(X_i))$ remains unchanged,
2. If $X_i \in \mathbf{S}$, then $P(X_i = x|pa(X_i)) = 1$, if x is the value that X_i was set to by the intervention $do(X_i = x)$; otherwise $P(X_i = x|pa(X_i)) = 0$.

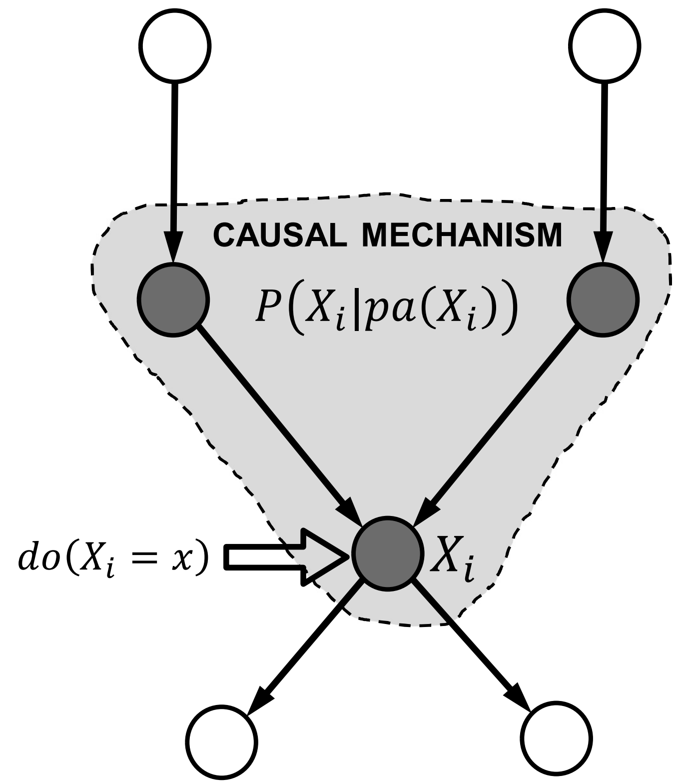

Figure 4.4(b)

We could write condition 2) below as follows:

- $P\big(X_i = x|pa(X_i)\big) = 1$ if x is consistent with the intervention
- $P\big(X_i = x|pa(X_i)\big) = 0$ otherwise.

In the future we say that, if $X_i \in \mathbf{S}$, a value x of X_i is consistent with the intervention on X_i, if x equals the value that X_i was set to in the intervention, i.e., $do(X_i = x)$.

The causal graph for interventional distributions is simply the same graph that was used for the observational joint distribution, but with all of the edges to the intervened node(s) removed.

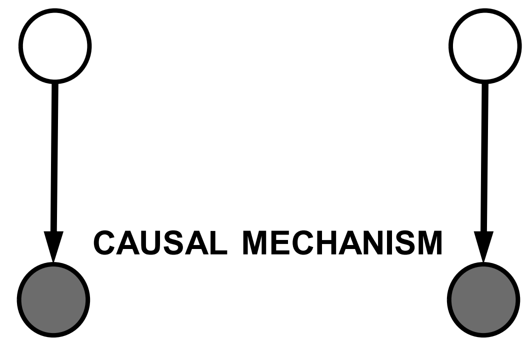

MODULARITY – INDEPENDENCE MECHANISM – INVARIANCE

If we intervene on a set of nodes/variables $\mathbf{S}$, setting them to constants, then for all $X_i \in \{X_1, X_2, \dots, X_n\}$, we have the following:

1. If $X_i \notin \mathbf{S}$, then $P\big(X_i = x|pa(X_i)\big)$ remains unchanged,
2. If $X_i \in \mathbf{S}$, then $P\big(X_i = x|pa(X_i)\big) = 1$, if x is the value that X_i was set to by the intervention $do(X_i = x)$; otherwise $P\big(X_i = x|pa(X_i)\big) = 0$.

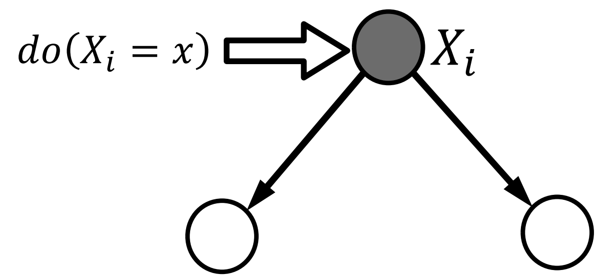

MANIPULATED GRAPH

Figure 4.5

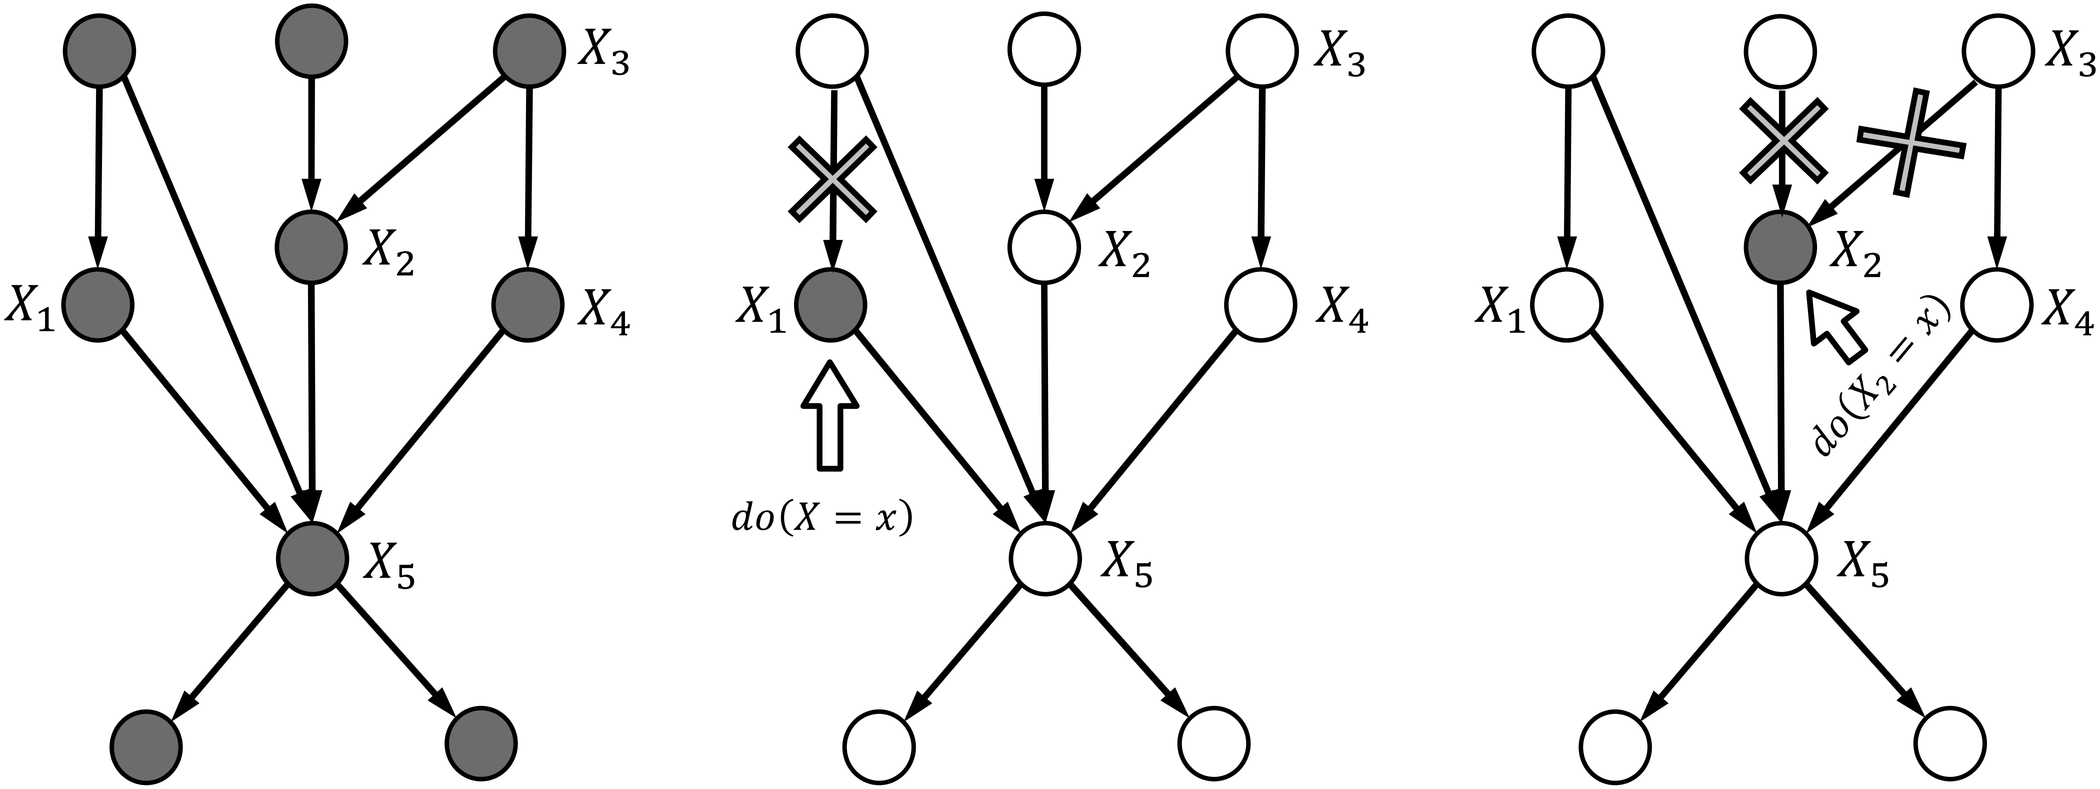

OBSERVATIONAL DISTRIBUTION

Figure 4.6(a)

INTERVENTIONAL DISTRIBUTION

Figure 4.6(b)

INTERVENTIONAL DISTRIBUTION

Figure 4.6(c)

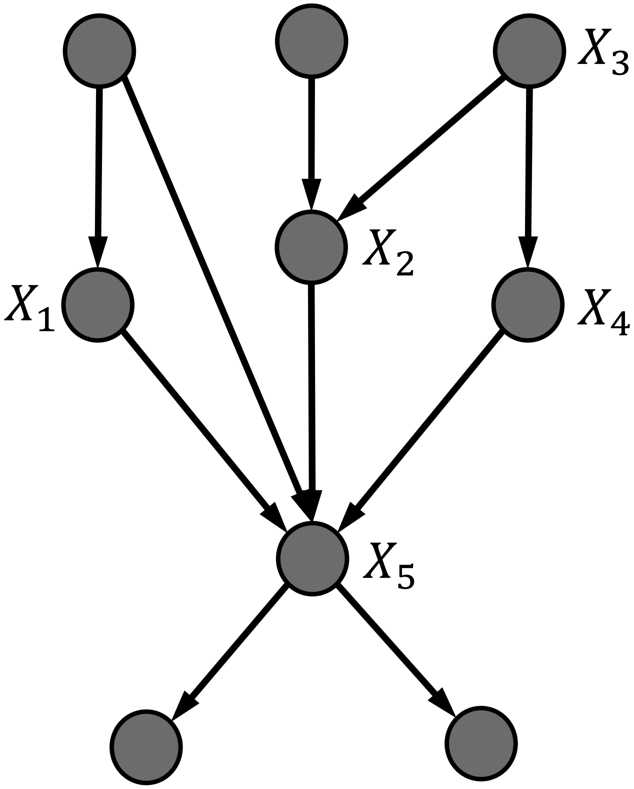

What would it mean for the **MODULARITY ASSUMPTION** to be violated?

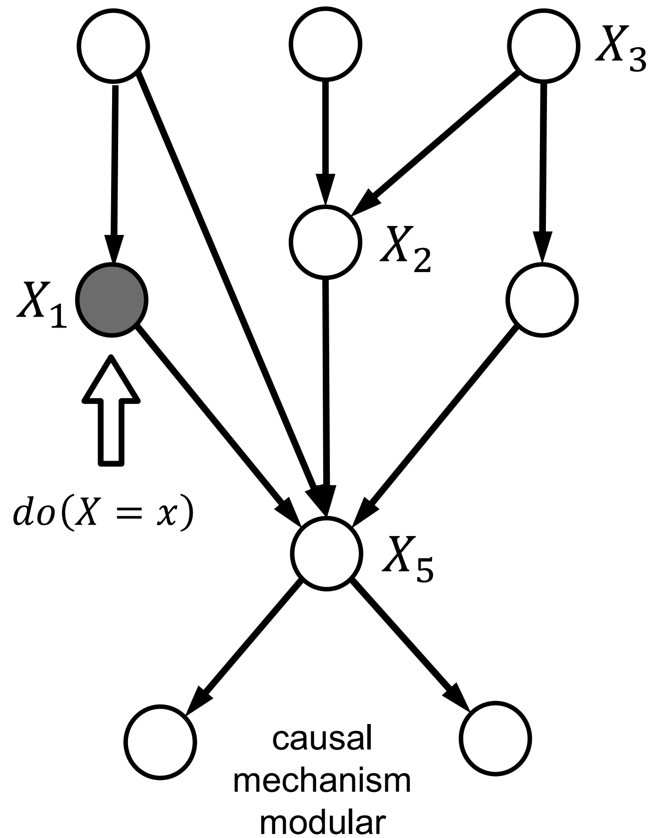

INTERVENTIONAL DISTRIBUTION

Figure 4.6(b)

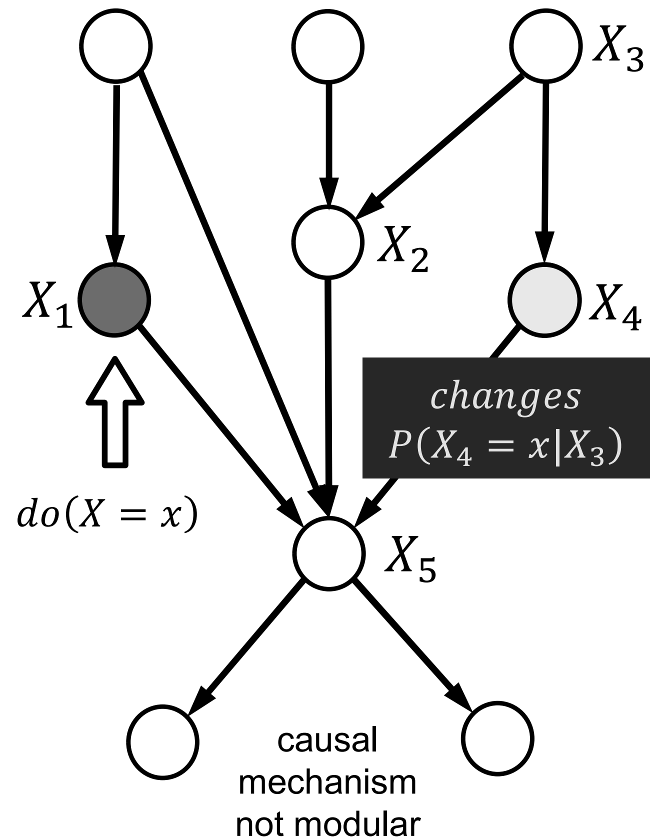

INTERVENTIONAL DISTRIBUTION

Figure 4.7

Using *do*-expressions and graph surgery, we can begin to untangle the causal relationships from the purely associative.

We now learn methods that can, astoundingly, tease out causal information from purely observational data, assuming of course that the graph constitutes a valid representation of reality.

It is worth noting here that we are making a tacit assumption that

> The **INTERVENTION** has "**NO SIDE EFFECTS**," that is, that assigning the value x for the variable X for an individual does not alter subsequent variables in a direct way.

For example,

- being "assigned" a drug might have a different effect on recovery than
- being forced to take the drug against one's religious objections.

When side effects are present, they need to be specified explicitly in the model.

The ice cream example represents an extreme case in which the association between X and Y was totally spurious from a causal perspective, because there was no causal path from X to Y.

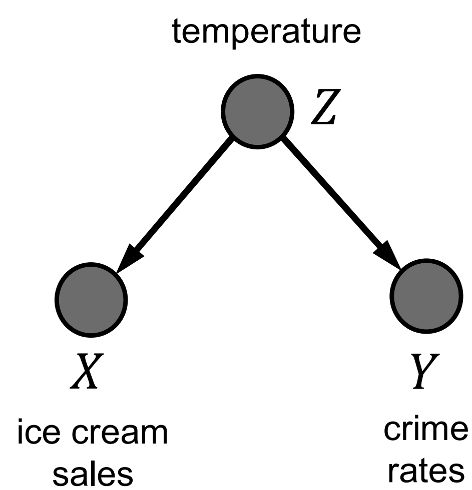

Figure 4.2

Let us come back to the **SIMPSON'S PARADOX**, where X stands for drug usage, Y stands for recovery, and Z stands for gender.

	Drug			No Drug		
	patients	recovered	% recovered	patients	recovered	% recovered
Men	87	81	93%	270	234	87%
Women	263	192	73%	80	55	69%
Combined data	350	273	78%	350	289	83%

To find out how effective the drug is in the population, we imagine a hypothetical intervention by which we administer the drug uniformly to the entire population $do(X=1)$ and compare the recovery rate to what would obtain under the complementary intervention, where we prevent everyone from using the drug $do(X=0)$.

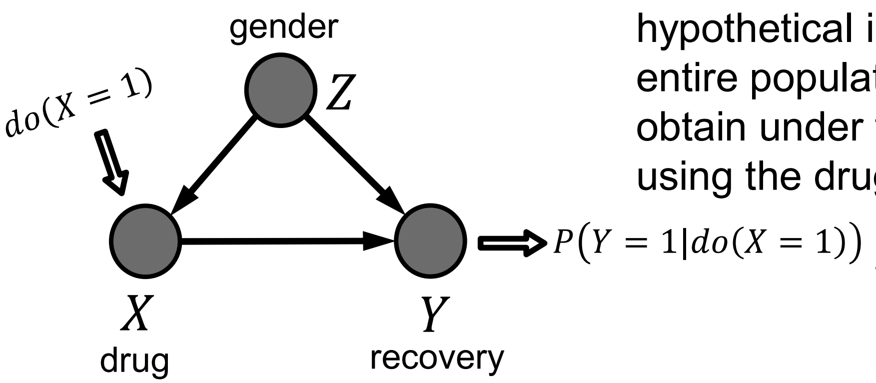

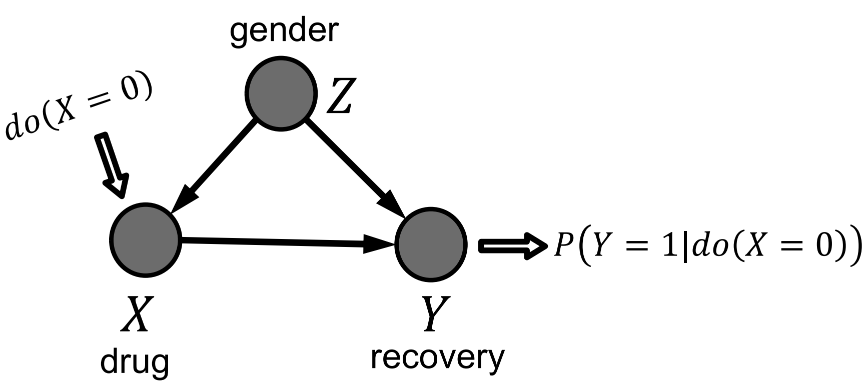

$$P(Y=1|do(X=1)) - P(Y=1|do(X=0))$$

AVERAGE TREATMENT EFFECT (ATE)

Let us come back to the **SIMPSON'S PARADOX**, where X stands for drug usage, Y stands for recovery, and Z stands for gender.

	Drug			No Drug		
	patients	recovered	% recovered	patients	recovered	% recovered
Men	87	81	93%	270	234	87%
Women	263	192	73%	80	55	69%
Combined data	350	273	78%	350	289	83%

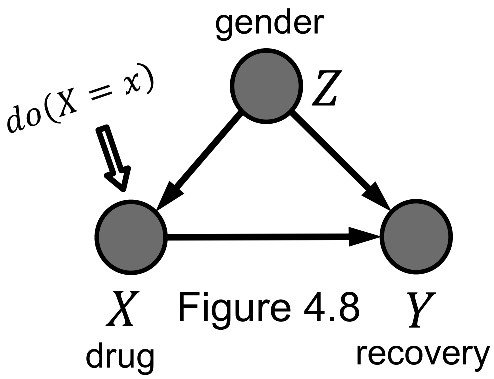

Figure 4.8

The data itself was not sufficient even for determining whether the effect of the drug was positive or negative.

But with the aid of the graph in Figure 4.8, we can compute the magnitude of the causal effect from the data (we generalize to more than two drugs and more than two outcomes).

To do so, we simulate the intervention in the form of a graph surgery on the **ORIGINAL MODEL** (Figure 4.8) just as we did in the ice cream example.

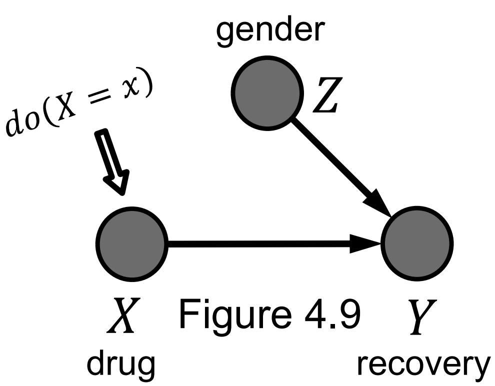

Figure 4.9

The intervention $do(X = x)$ brings to the **MANIPULATED MODEL** in Figure 4.9, which in turns allows us to write the following equality:

$$P\big(Y = y \mid do(X = x)\big) = P_m(Y = y \mid X = x)$$

However, to draw the above conclusion we need to assume **MODULARITY – INDEPENDENCE MECHANISM – INVARIANCE.**

Let us come back to the **SIMPSON'S PARADOX**, where X stands for drug usage, Y stands for recovery, and Z stands for gender.

	Drug			No Drug		
	patients	recovered	% recovered	patients	recovered	% recovered
Men	87	81	93%	270	234	87%
Women	263	192	73%	80	55	69%
Combined data	350	273	78%	350	289	83%

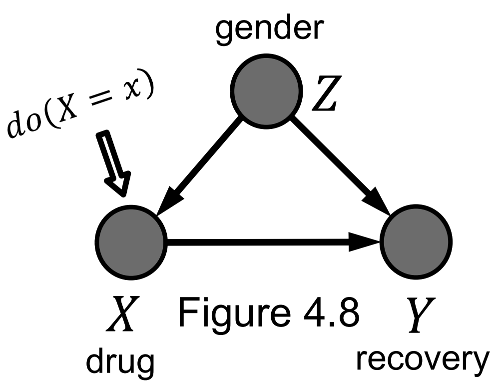

Figure 4.8

In other terms to computing the causal effect we assume that the manipulated probability (manipulated model or post-intervention model in Figure 4.9), shares two essential properties with the original probability, that prevails in the original model or pre-intervention model of Figure 4.8, i.e.,:

$$P\big(Z|do(X = x)\big) = P_m(Z|x) = P(Z)$$

The marginal probability $P(Z)$ is invariant under the intervention, because the process determining Z is not affected by removing the arrow from Z to X. (proportions of males and females remain the same, before and after the intervention).

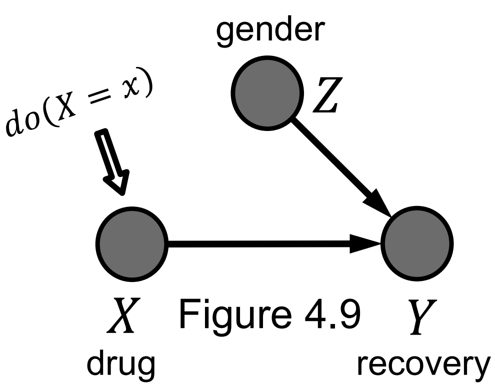

Figure 4.9

$$P(Y|do(X = x), Z) = P_m(Y|x, Z) = P(Y|x, Z)$$

The conditional probability $P(Y|x, Z)$ is invariant, because the process by which Y responds to X and Z remains the same, regardless of whether X changes spontaneously or by deliberate manipulation.

We know that the following equalities hold:

pre-intervention distribution P

$$P_m(Z=z|X=x)=P_m(Z=z)=P(Z=z)$$

$$P_m(Y=y|X=x,Z=z)=P(Y=y|X=x,Z=z)$$

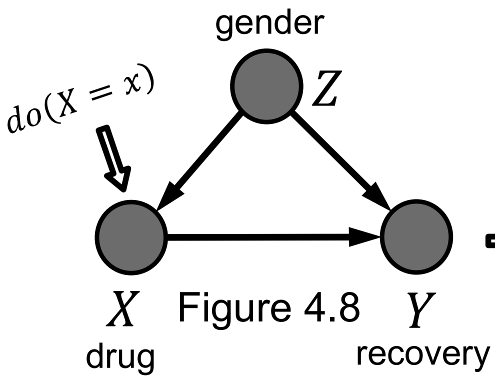

Figure 4.8

$$\Longrightarrow P\big(Y=y|do(X=x)\big)=P_m(Y=y|X=x) \qquad \text{(by definition)}$$

$$=\sum_z P_m(Y=y|X=x,Z=z)\,P_m(Z=z|X=x)$$

$$=\sum_z P_m(Y=y|X=x,Z=z)\,P_m(Z=z)$$

$$=\sum_z P(Y=y|X=x,Z=z)\,P(Z=z) \qquad \text{(modularity)}$$

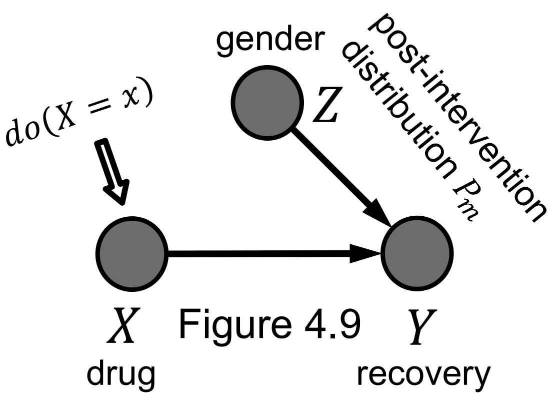

Figure 4.9

We know that the following equalities hold:

pre-intervention distribution P

$$P_m(Z = z|X = x) = P_m(Z = z) = P(Z = z)$$

$$P_m(Y = y|X = x, Z = z) = P(Y = y|X = x, Z = z)$$

Computes the association between X and Y for each value z of Z, then averages over those values.

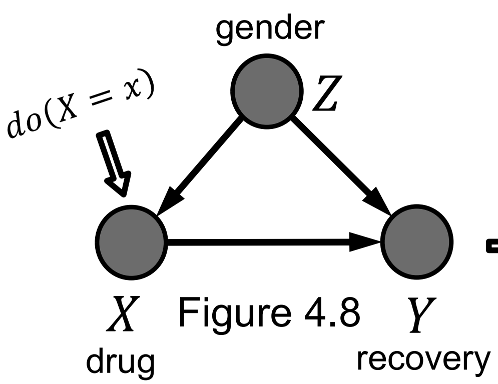

Figure 4.8

ADJUSTMENT FORMULA

$$P\big(Y = y|do(X = x)\big) = \sum_z P(Y = y|X = x, Z = z)\, P(Z = z)$$

ADJUSTING FOR *Z*
or
CONTROLLING FOR *Z*

It can be estimated directly from the data, since it consists only of conditional probabilities, i.e., by the **PRE-INTERVENTION OR OBSERVATIONAL DISTRIBUTION *P***.

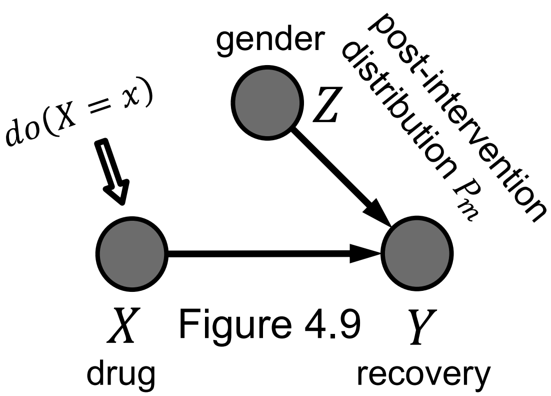

Figure 4.9

THIS CAUSAL EXPRESSION IS IDENTIFIABLE

CAUSAL ESTIMAND

$$P\big(Y = y|do(X = x)\big)$$

IDENTIFICATION →

STATISTICAL ESTIMAND

$$\sum_z P(Y = y|X = x, Z = z)\, P(Z = z)$$

We now show that the Adjustment Formula works, while using the Simpson's paradox.

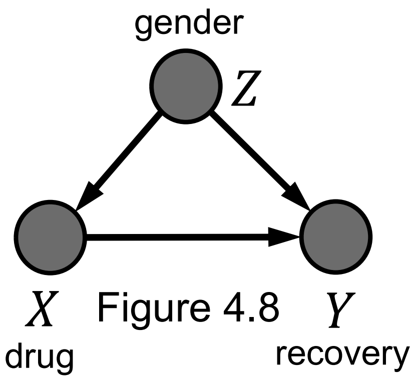

Figure 4.8

To demonstrate the working of the adjustment formula, let us apply it numerically to Simpson's story, with

- $X = 1$ standing for the patient taking the drug
- $Z = 1$ standing for the patient being male
- $Y = 1$ standing for the patient recovering

ADJUSTMENT FORMULA

$$P\big(Y = y | do(X = x)\big) = \sum_z P(Y = y | X = x, Z = z)\, P(Z = z)$$

$$P\big(Y = 1 | do(X = 1)\big) = P(Y = 1 | X = 1, Z = 1)\, P(Z = 1) + P(Y = 1 | X = 1, Z = 0)\, P(Z = 0)$$

We now show that the Adjustment Formula works, while using the Simpson's paradox.

	Drug			No Drug		
	patients	recovered	% recovered	patients	recovered	% recovered
Men	87	81	93%	270	234	87%
Women	263	192	73%	80	55	69%
Combined data	350	273	78%	350	289	83%

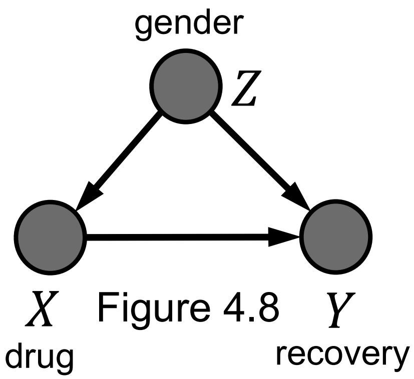

Figure 4.8

ADJUSTMENT FORMULA

$$P\big(Y = y | do(X = x)\big) = \sum_z P(Y = y | X = x, Z = z)\, P(Z = z)$$

$$P\big(Y = 1 | do(X = 1)\big) = P(Y = 1 | X = 1, Z = 1)\, P(Z = 1) + P(Y = 1 | X = 1, Z = 0)\, P(Z = 0)$$

$$P\big(Y = 1 | do(X = 1)\big) = 0.93 \times \frac{(87 + 270)}{700} + 0.73 \times \frac{(263 + 80)}{700} = 0.832$$

$$P\big(Y = 1 | do(X = 0)\big) = 0.87 \times \frac{(87 + 270)}{700} + 0.69 \times \frac{(263 + 80)}{700} = 0.7818$$

A clear positive advantage to drug-taking

$$\text{ATE} = P\big(Y = 1 | do(X = 1)\big) - P\big(Y = 1 | do(X = 0)\big) = 0.832 - 0.7818 = 0.0502$$

We see that the adjustment formula instructs us to

- condition on gender Z,
- find the benefit of the drug separately for males and females,
- average the result using the percentage of males and females in the population.

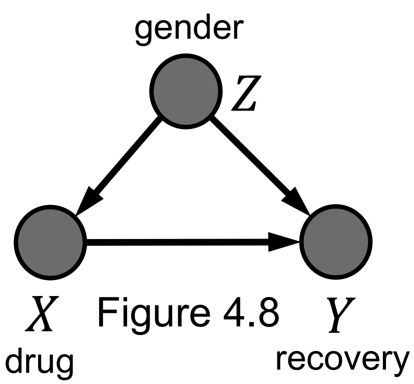

Figure 4.8

ADJUSTMENT FORMULA

$$P\big(Y = y|do(X = x)\big) = \sum_{z} P(Y = y|X = x, Z = z)\; P(Z = z)$$

probability to recover for male drug-takers

average using percentage male

probability to recover for female drug-takers

average using percentage female

$$P\big(Y = 1|do(X = 1)\big) = P(Y = 1|X = 1, Z = 1)\, P(Z = 1) + P(Y = 1|X = 1, Z = 0)\, P(Z = 0)$$

condition on gender (male)

condition on gender (female)

It also thus instructs us to ignore the aggregated population data

$P(Y = 1|X = 1)$ $P(Y = 1|X = 0)$

from which we might (falsely) conclude that the drug has a negative effect overall.

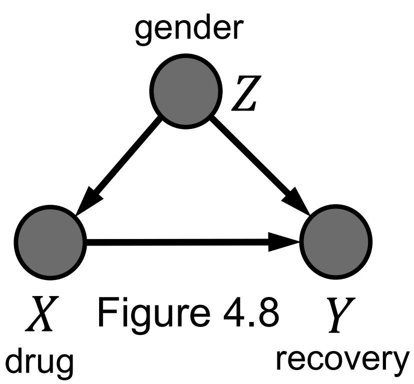

Figure 4.8

$$\text{ATE} = P(Y = 1|X = 1) - P(Y = 1|X = 0) = 0.78 - 0.83 = -0.05$$

⇨ The drug appears to be harmful

ADJUSTMENT FORMULA

$$P\big(Y = y|do(X = x)\big) = \sum_z P(Y = y|X = x, Z = z)\, P(Z = z)$$

	Drug			No Drug		
	patients	recovered	% recovered	patients	recovered	% recovered
Men	87	81	93%	270	234	87%
Women	263	192	73%	80	55	69%
Combined data	350	273	78%	350	289	83%

These simple examples might give us the impression that whenever we face the dilemma of whether to condition on a third variable Z, the adjustment formula prefers the Z-specific analysis over the nonspecific analysis.

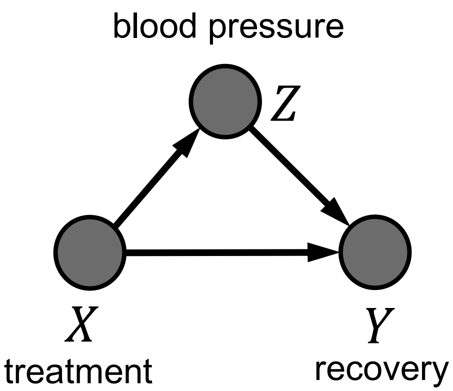

Figure 4.10

But what about the blood pressure example of Simpson's paradox?

	No Drug			Drug		
	patients	recovered	% recovered	patients	recovered	% recovered
Low BP	87	81	93%	270	234	87%
High BP	263	192	73%	80	55	69%
Combined data	350	273	78%	350	289	83%

The more sensible method would be not to condition on blood pressure, but to examine the unconditional population table directly.

$$P\big(Y = 1 | do(X = 1)\big) = ?$$

How would the adjustment formula cope with situations like that?

These simple examples might give us the impression that whenever we face the dilemma of whether to condition on a third variable Z, the adjustment formula prefers the Z-specific analysis over the nonspecific analysis.

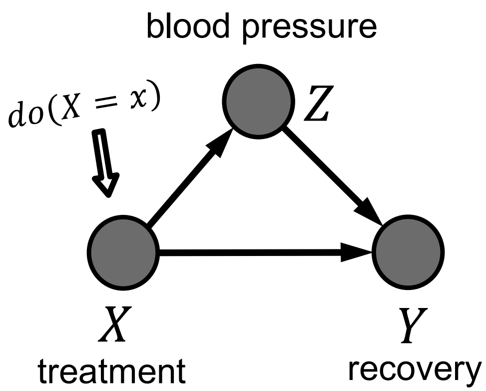

Figure 4.10

But what about the blood pressure example of Simpson's paradox?

	No Drug			Drug		
	patients	recovered	% recovered	patients	recovered	% recovered
Low BP	87	81	93%	270	234	87%
High BP	263	192	73%	80	55	69%
Combined data	350	273	78%	350	289	83%

- We simulate an intervention and then examine the adjustment formula that emanates from the simulated intervention.
- In graphical models, an intervention is simulated by severing all arrows that enter the manipulated variable X.

X has no entering edges

↓

no surgery needed ⟶ $P(Y = 1 | do(X = 1)) = P(Y = 1 | X = 1)$

pre-intervention distribution P is the same as the post-intervention distribution P_m

We are now in a position to understand what variable Z, or set of variables $\mathbf{Z}$, can legitimately be included in the adjustment formula.

The intervention procedure, which led to the adjustment formula, dictates that $\mathbf{Z}$ should coincide with the parents $pa(X)$ of X, because it is the influence of these parents that we neutralize when we fix X by external manipulation $do(X = x)$.

We can therefore write a general adjustment formula and summarize it in a rule:

THE CAUSAL EFFECT RULE

Given a graph $\mathcal{G}$ in which a set of variables $pa(X)$ are designated as the parents of X, the causal effect of X on Y is given by

$$P\big(Y = y | do(X = x)\big) = \sum_{\boldsymbol{u}} P(Y = y | X = x, pa(X) = \boldsymbol{u})\, P(pa(X) = \boldsymbol{u})$$

where $\boldsymbol{u}$ ranges over all the combinations of values that the variables in $pa(X)$ can take.

$$P\big(Y = y | do(X = x)\big) = \sum_{\boldsymbol{u}} \frac{P(Y = y, X = x, pa(X) = \boldsymbol{u})}{\underbrace{P(X = x | pa(X) = \boldsymbol{u})}}$$

PROPENSITY SCORE: displays the role played by the parents $pa(X)$ of X in predicting the results of interventions, the advantages of expressing $P\big(y|do(x)\big)$ in this form will be not discussed here.

THE CAUSAL EFFECT RULE

Given a graph $\mathcal{G}$ in which a set of variables $pa(X)$ are designated as the parents of X, the causal effect of X on Y is given by

$$P\big(Y = y | do(X = x)\big) = \sum_{\boldsymbol{u}} \boxed{P(Y = y | X = x, pa(X) = \boldsymbol{u})\, P(pa(X) = \boldsymbol{u})}$$

where $\boldsymbol{u}$ ranges over all the combinations of values that the variables in $pa(X)$ can take.

If we multiply and divide the summand by the probability $P(X = x | pa(X) = \boldsymbol{u})$, we get a more convenient form:

$$P\big(Y = y | do(X = x)\big) = \sum_{\boldsymbol{u}} \frac{P(Y = y, X = x, pa(X) = \boldsymbol{u})}{\underbrace{P(X = x | pa(X) = \boldsymbol{u})}}$$

PROPENSITY SCORE: displays the role played by the parents $pa(X)$ of X in predicting the results of interventions, the advantages of expressing $P\big(y|do(x)\big)$ in this form will be not discussed here.

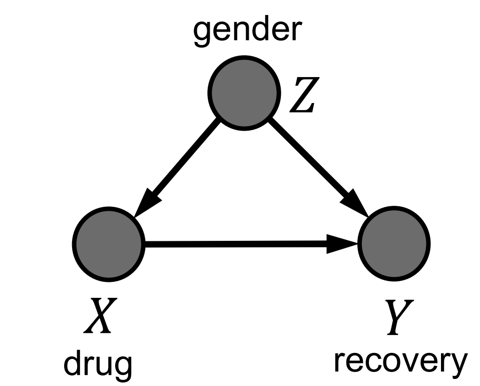

Figure 4.8

We can appreciate now what role the causal graph plays in resolving Simpson's paradox, and, more generally, what aspects of the graph allow us to predict causal effects from purely observational data.

We need the graph in order to determine the identity of X's parents $pa(X)$—the set of factors that, under non-experimental conditions, would be sufficient for determining the value of X, or the probability of that value.

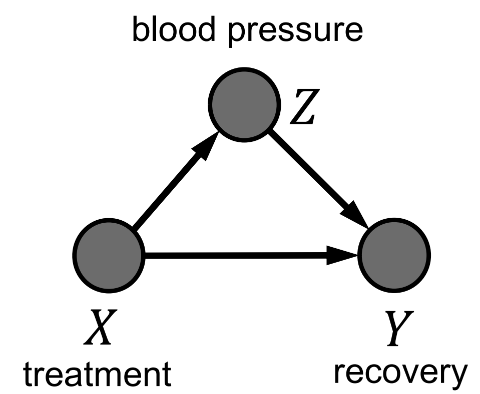

Figure 4.10

Using graphs and their underlying assumptions, we were able to identify causal relationships in purely observational data.

But, from this discussion, readers may be tempted to conclude that the role of graphs is fairly limited; once we identify the parents $pa(X)$ of X, the rest of the graph can be discarded, and the causal effect can be evaluated mechanically from the adjustment formula.

The next part of the lecture shows that things may not be so simple.

PART III

TRUNCATED FACTORIZATION AND BACKDOOR ADJUSTMENT

ADJUSTMENT FORMULA

$$P\big(Y = y|do(X = x)\big) = \sum_z P(Y = y|X = x, Z = z)\, P(Z = z)$$

In deriving the adjustment formula, we assumed

- an intervention on a single variable X,
- whose parents were disconnected,

so as to simulate the absence of their influence after intervention.

However, social/medical policies occasionally involve **MULTIPLE INTERVENTIONS**, such as those that dictate the value of several variables simultaneously, or those that control a variable over time.

Then, it is useful to start from the:

BAYESIAN NETWORK FACTORIZATION

Given a probability distribution P and a DAG $\mathcal{G}$, P factorizes according to $\mathcal{G}$ if

$$P(X_1, X_2, \dots, X_n) = \prod_{i=1}^{n} P\big(X_i|pa(X_i)\big)$$

PRE-INTERVENTION DISTRIBUTION

$$P(x, y, z) = P(z)\, P(x|z)\, P(y|x, z)$$

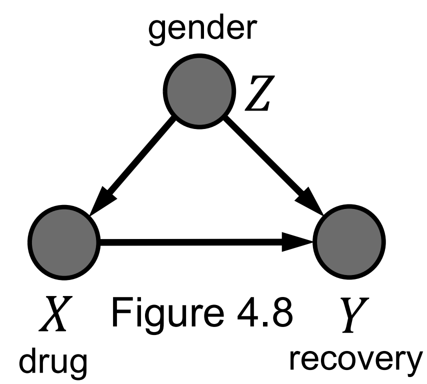

Figure 4.8

POST-INTERVENTION DISTRIBUTION

$$P\big(z, y|do(x)\big) = P_m(z)P_m(x|z)P_m(y|x, z) = P(z)\, P(y|x, z)$$

$$P_m(x|z) = 1$$

$$P\big(y|do(x)\big) = \sum_z P\big(z, y|do(x)\big)$$

$$= \sum_z P(z)P(y|x, z) \neq P(y|x)$$

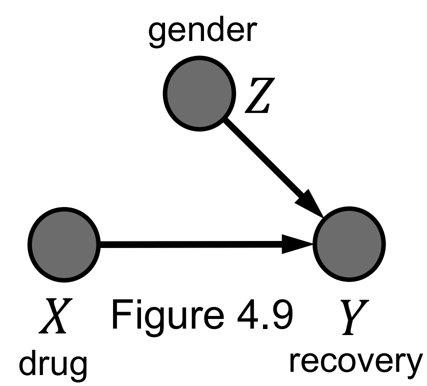

Figure 4.9

ADJUSTMENT FORMULA

$$P\big(Y = y|do(X = x)\big) = \sum_z P(Y = y|X = x, Z = z)\ P(Z = z)$$

However, social and medical policies occasionally involve multiple interventions, such as those that dictate the value of several variables simultaneously, or those that control a variable over time.

Then, it is useful to start from the:

BAYESIAN NETWORK FACTORIZATION

Given a probability distribution P and a DAG $\mathcal{G}$, P factorizes according to $\mathcal{G}$ if

$$P(X_1, X_2, \dots, X_n) = \prod_{i=1}^{n} P\big(X_i|pa(X_i)\big)$$

In deriving the adjustment formula, we assumed

- an intervention on a single variable X,
- whose parents were disconnected,

so as to simulate the absence of their influence after intervention.

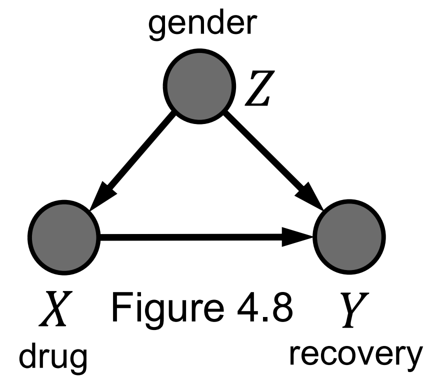

Figure 4.8

If $P(z) = P(z|x)$ then, we would have

$$\sum_z P(z)P(y|x,z)$$

$$= \sum_z P(z|x)P(y|x,z)$$

$$P\big(y|do(x)\big) = \sum_z P(y,z|x) = P(y|x)$$

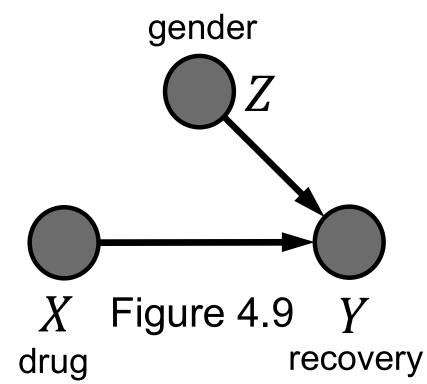

Figure 4.9

$$P\big(y|do(x)\big) = \sum_z P\big(z,y|do(x)\big) \neq P(y|x)$$

$$= \sum_z P(z)P(y|x,z)$$

ADJUSTMENT FORMULA

$$P(Y=y|do(X=x)) = \sum_z P(Y=y|X=x, Z=z)\, P(Z=z)$$

In deriving the adjustment formula, we assumed

- an intervention on a single variable X,
- whose parents were disconnected,

so as to simulate the absence of their influence after intervention.

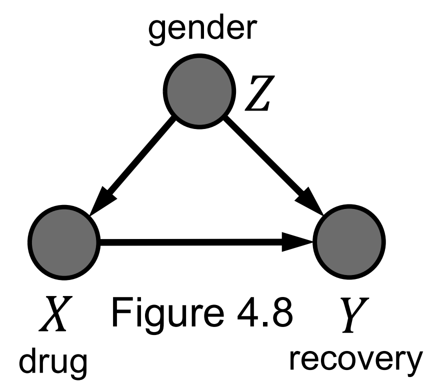

Figure 4.8

$$\text{ATE} = \mathbb{E}[Y(1) - Y(0)] = \sum_y y\, P(y|do(X=1)) - \sum_y y\, P(y|do(X=0))$$

BAYESIAN NETWORK FACTORIZATION

Given a probability distribution P and a DAG $\mathcal{G}$, P factorizes according to $\mathcal{G}$ if

$$P(X_1, X_2, \dots, X_n) = \prod_{i=1}^{n} P(X_i|pa(X_i))$$

$$P(y|do(x)) = \sum_z P(z, y|do(x)) \neq P(y|x)$$

$$= \sum_z P(z)P(y|x, z)$$

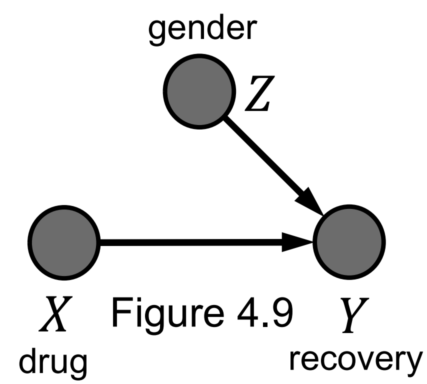

Figure 4.9

ADJUSTMENT FORMULA

$$P\big(Y = y|do(X = x)\big) = \sum_{z} P(Y = y|X = x, Z = z)\ P(Z = z)$$

The previous consideration also allows us to generalize the **ADJUSTMENT FORMULA** to **MULTIPLE INTERVENTIONS**, that is, interventions that fix the values of a set of variables $\mathbf{S}$ to constants $\boldsymbol{s}$.

We simply write down the **FACTORIZATION** of the **PRE-INTERVENTION DISTRIBUTION** and strike out all factors that correspond to variables in the **INTERVENTION SET** $\mathbf{S}$.

TRUNCATED FACTORIZATION – G-FORMULA

We assume that P and $\mathcal{G}$ satisfy the Markov assumption and modularity. Given, a set of intervention nodes $\mathbf{S}$ (intervention set), if x_i is consistent with the intervention $\mathbf{S} = \boldsymbol{s}$, then

$$P\big(x_1, x_2, \dots, x_n|do(\mathbf{S} = \boldsymbol{s})\big) = \prod_{X_i \notin \mathbf{S}}^{n} P\big(X_i = x_i|pa(X_i)\big)$$

otherwise $P\big(x_1, x_2, \dots, x_n|do(\mathbf{S} = \boldsymbol{s})\big) = 0$.

BAYESIAN NETWORK FACTORIZATION

Given a probability distribution P and a DAG $\mathcal{G}$, P factorizes according to $\mathcal{G}$ if

$$P(X_1, X_2, \dots, X_n) = \prod_{i=1}^{n} P\big(X_i|pa(X_i)\big)$$

PRE-INTERVENTION DISTRIBUTION

$$P\big(z_1, z_2, w, y \,|\, \underbrace{do(X = x, Z_3 = z_3)}_{do(\mathbf{S} = \boldsymbol{s})}\big) = \overbrace{P(z_1)\, P(z_2)\, P(x|z_1, z_3)\, P(z_3|z_1, z_2)\, P(w|x)\, P(y|w, z_2, z_3)}^{\text{PRE-INTERVENTION DISTRIBUTION}}$$

$$= P(z_1)\, P(z_2) \quad 1 \quad 1 \quad P(w|x)\, P(y|w, z_2, z_3)$$

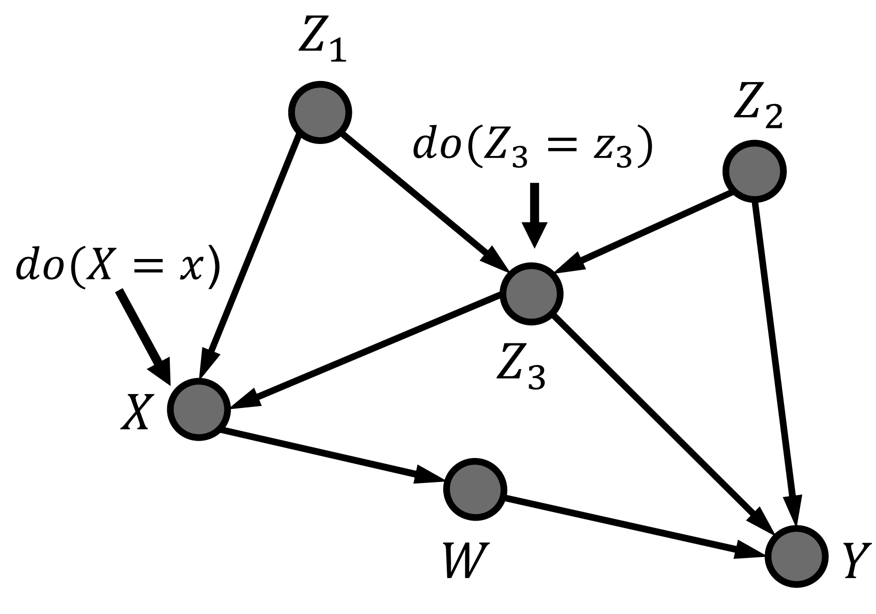

Figure 4.11

TRUNCATED FACTORIZATION – G-FORMULA

We assume that P and $\mathcal{G}$ satisfy the Markov assumption and modularity. Given, a set of intervention nodes $\mathbf{S}$ (intervention set), if x_i is consistent with the intervention $\mathbf{S} = \boldsymbol{s}$, then

$$P\big(x_1, x_2, \dots, x_n | do(\mathbf{S} = \boldsymbol{s})\big) = \prod_{X_i \notin \mathbf{S}}^{n} P\big(X_i = x_i | pa(X_i)\big)$$

otherwise $P\big(x_1, x_2, \dots, x_n | do(\mathbf{S} = \boldsymbol{s})\big) = 0$.

PRE-INTERVENTION DISTRIBUTION

$$P\big(z_1, z_2, w, y \,|\, \underbrace{do(X = x, Z_3 = z_3)}_{do(\mathbf{S}=\boldsymbol{s})}\big) = \overbrace{P(z_1)\, P(z_2)\, P(x|z_1, z_3)\, P(z_3|z_1, z_2)\, P(w|x)\, P(y|w, z_2, z_3)}$$

$$= P(z_1)\, P(z_2) \quad 1 \quad 1 \quad P(w|x)\, P(y|w, z_2, z_3)$$

POST-INTERVENTION DISTRIBUTION

$$P\big(z_1, z_2, w, y | do(X = x, Z_3 = z_3)\big) = P(z_1)\, P(z_2)\, P(w|x)\, P(y|w, z_2, z_3)$$

TRUNCATED FACTORIZATION – G-FORMULA

We assume that P and $\mathcal{G}$ satisfy the Markov assumption and modularity. Given, a set of intervention nodes $\mathbf{S}$ (intervention set), if x_i is consistent with the intervention $\mathbf{S} = \boldsymbol{s}$, then

$$P\big(x_1, x_2, \ldots, x_n | do(\mathbf{S} = \boldsymbol{s})\big) = \prod_{X_i \notin \mathbf{S}}^{n} P\big(X_i = x_i | pa(X_i)\big)$$

otherwise $P\big(x_1, x_2, \ldots, x_n | do(\mathbf{S} = \boldsymbol{s})\big) = 0$.

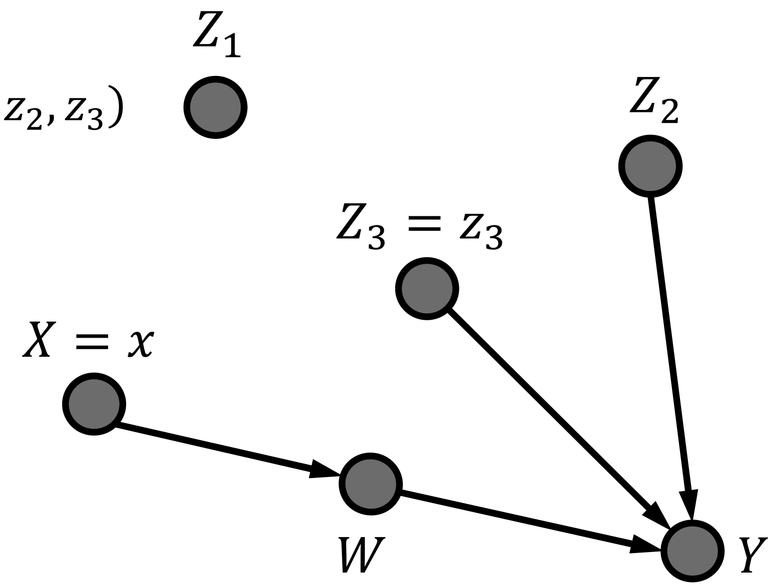

Figure 4.12

It is interesting to note that combining

$$P(x, y, z) = P(z)\, P(x|z)\, P(y|x, z) \qquad \text{(pre-intervention)}$$

and

$$P\big(z, y|do(x)\big) = P_m(z)\, P_m(y|x, z) = P(z)\, P(y|x, z) \qquad \text{(post-intervention)}$$

we get a simple relation between the pre-and post-intervention distributions:

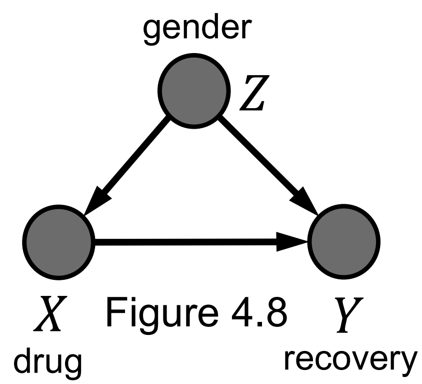

Figure 4.8

$$P\big(z, y|do(x)\big) = \frac{P(x, y, z)}{P(x|z)}$$

It tells us that the conditional probability $P(x|z)$ is all we need to know in order to predict the effect of an intervention $do(x)$ from non-experimental data governed by the distribution $P(x, y, z)$.

TRUNCATED FACTORIZATION – G-FORMULA

We assume that P and $\mathcal{G}$ satisfy the Markov assumption and modularity. Given, a set of intervention nodes $\mathbf{S}$ (intervention set), if x_i is consistent with the intervention $\mathbf{S} = \boldsymbol{s}$, then

$$P\big(x_1, x_2, \dots, x_n|do(\mathbf{S} = \boldsymbol{s})\big) = \prod_{X_i \notin \mathbf{S}}^{n} P\big(X_i = x_i|pa(X_i)\big)$$

otherwise $P\big(x_1, x_2, \dots, x_n|do(\mathbf{S} = \boldsymbol{s})\big) = 0$.

THE CAUSAL EFFECT RULE

Given a graph $\mathcal{G}$ in which a set of variables $pa(X)$ are designated as the parents of X, the causal effect of X on Y is given by

$$P\big(Y = y|do(X = x)\big) = \sum_{\boldsymbol{u}} P(Y = y|X = x, pa(X) = \boldsymbol{u})\, P(pa(X) = \boldsymbol{u})$$

where $\boldsymbol{u}$ ranges over all the combinations of values that the variables in $pa(X)$ can take.

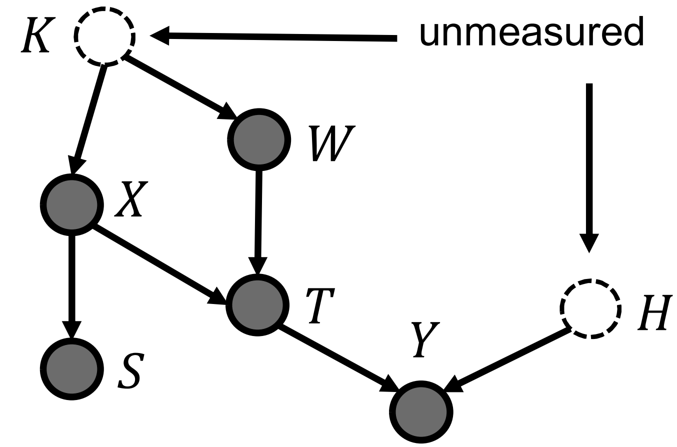

Figure 4.13

We came to the conclusion that, given a variable X, we should adjust for it's parents $pa(X)$, when trying to determine the effect of X on another variable Y.

But often, we know, or believe, that the variables have **UNMEASURED PARENTS (LATENT)** that, though represented in the graph, may be inaccessible for measurement.

In those cases, we need to find an alternative set of variables to adjust for.

Under what conditions does a causal story permit us to compute the causal effect of one variable on another, from data obtained by passive observations, with no interventions?

Since we have decided to represent causal stories with graphs, the question becomes a graph-theoretical problem:

> Under what conditions, is the structure of the causal graph sufficient for computing a causal effect from a given data set?

The answer to that question is long enough—and important enough—that we will spend the rest of the lecture addressing it.

But one of the most important tools we use to determine whether we can compute a causal effect is a simple test called the **BACKDOOR CRITERION.**

Using it, we can determine, for any two variables X and Y in a causal model represented by a DAG $\mathcal{G}$, which set of variables $\mathbf{S}$ in that model should be conditioned on when searching for the causal relationship between X and Y.

THE BACKDOOR CRITERION

Given an ordered pair of variables (X, Y) in a DAG $\mathcal{G}$, a set of variables $\mathbf{S}$ satisfies the backdoor criterion relative to (X, Y) if no node in $\mathbf{S}$ is a descendant of X, and $\mathbf{S}$ blocks every path between X and Y that contains an arrow into X.

If a set of variables $\mathbf{S}$ satisfies the **BACKDOOR CRITERION** for X and Y, then the causal effect of X on Y is given by the formula

THE BACKDOOR ADJUSTMENT FORMULA

$$P(Y = y | do(X = x)) = \sum_{\boldsymbol{s}} P(Y = y | X = x, \mathbf{S} = \boldsymbol{s})\, P(\mathbf{S} = \boldsymbol{s})$$

just as when we adjust for $pa(X)$.

(Note that $pa(X)$ always satisfies the backdoor criterion)

The logic behind the **BACKDOOR CRITERION** is fairly straightforward.

In general, we would like to condition on a set of nodes $\mathbf{S}$ (**CONDITIONING SET**) such that we:

1. block all spurious paths between X and Y
2. leave all directed paths from X to Y unperturbed
3. create no new spurious paths

1. block all spurious paths between *X* and *Y*.

We want the conditioning set $\mathbf{S}$ to block any **BACKDOOR PATH** in which one end has an arrow into X, because such paths may make X and Y dependent, but are obviously not transmitting causal influences from X, and if we do not block them, they will confound the effect that X has on Y.

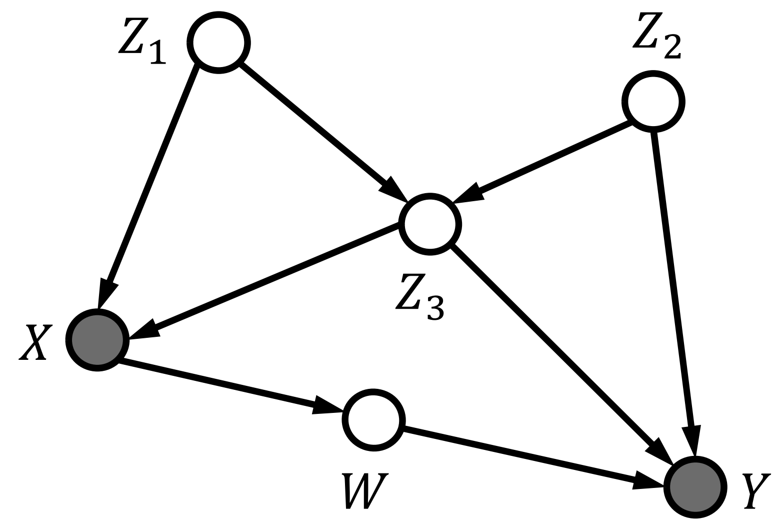

Figure 4.14

THE BACKDOOR CRITERION

Given an ordered pair of variables (X, Y) in a DAG $\mathcal{G}$, a set of variables $\mathbf{S}$ satisfies the backdoor criterion relative to (X, Y) if no node in $\mathbf{S}$ is a descendant of X, and $\mathbf{S}$ blocks every path between X and Y that contains an arrow into X.

The logic behind the **BACKDOOR CRITERION** is fairly straightforward.

In general, we would like to condition on a set of nodes $\mathbf{S}$ (**CONDITIONING SET**) such that we:

1. block all spurious paths between X and Y
2. leave all directed paths from X to Y unperturbed
3. create no new spurious paths

1. block all spurious paths between *X* and *Y*.

We want the conditioning set $\mathbf{S}$ to block any **BACKDOOR PATH** in which one end has an arrow into X, because such paths may make X and Y dependent, but are obviously not transmitting causal influences from X, and if we do not block them, they will confound the effect that X has on Y.

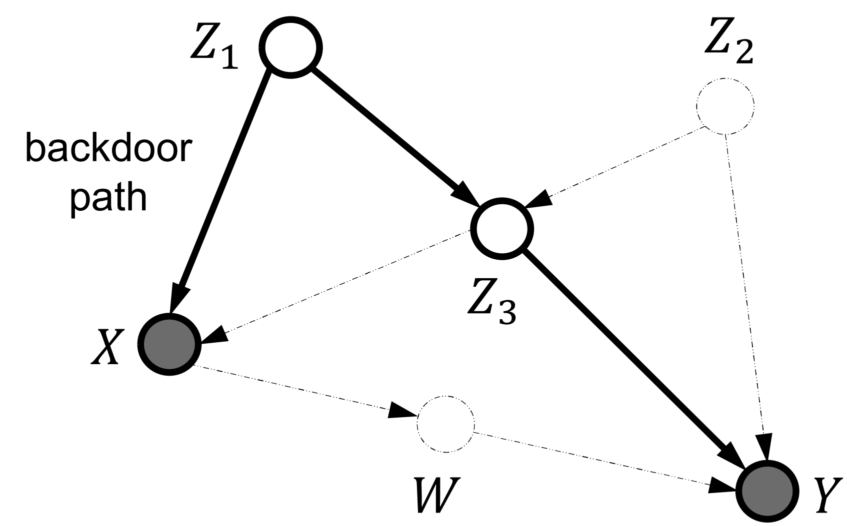

Figure 4.15

THE BACKDOOR CRITERION

Given an ordered pair of variables (X, Y) in a DAG $\mathcal{G}$, a set of variables $\mathbf{S}$ satisfies the backdoor criterion relative to (X, Y) if no node in $\mathbf{S}$ is a descendant of X, and $\mathbf{S}$ blocks every path between X and Y that contains an arrow into X.

The logic behind the **BACKDOOR CRITERION** is fairly straightforward.

In general, we would like to condition on a set of nodes $\mathbf{S}$ (**CONDITIONING SET**) such that we:

1. block all spurious paths between X and Y
2. leave all directed paths from X to Y unperturbed
3. create no new spurious paths

1. block all spurious paths between *X* and *Y*.

We want the conditioning set $\mathbf{S}$ to block any **BACKDOOR PATH** in which one end has an arrow into X, because such paths may make X and Y dependent, but are obviously not transmitting causal influences from X, and if we do not block them, they will confound the effect that X has on Y.

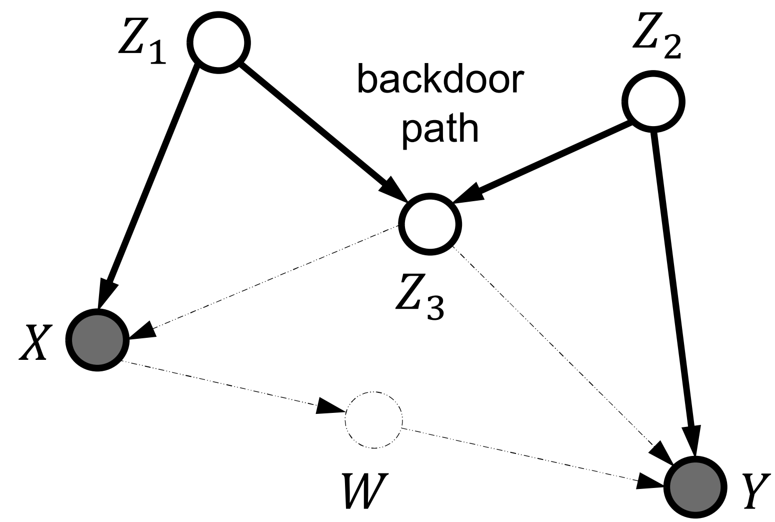

Figure 4.16

THE BACKDOOR CRITERION

Given an ordered pair of variables (X, Y) in a DAG $\mathcal{G}$, a set of variables $\mathbf{S}$ satisfies the backdoor criterion relative to (X, Y) if no node in $\mathbf{S}$ is a descendant of X, and $\mathbf{S}$ blocks every path between X and Y that contains an arrow into X.

The logic behind the **BACKDOOR CRITERION** is fairly straightforward.

In general, we would like to condition on a set of nodes $\mathbf{S}$ (**CONDITIONING SET**) such that we:

1. block all spurious paths between X and Y
2. leave all directed paths from X to Y unperturbed
3. create no new spurious paths

1. block all spurious paths between *X* and *Y*.

We want the conditioning set $\mathbf{S}$ to block any **BACKDOOR PATH** in which one end has an arrow into X, because such paths may make X and Y dependent, but are obviously not transmitting causal influences from X, and if we do not block them, they will confound the effect that X has on Y.

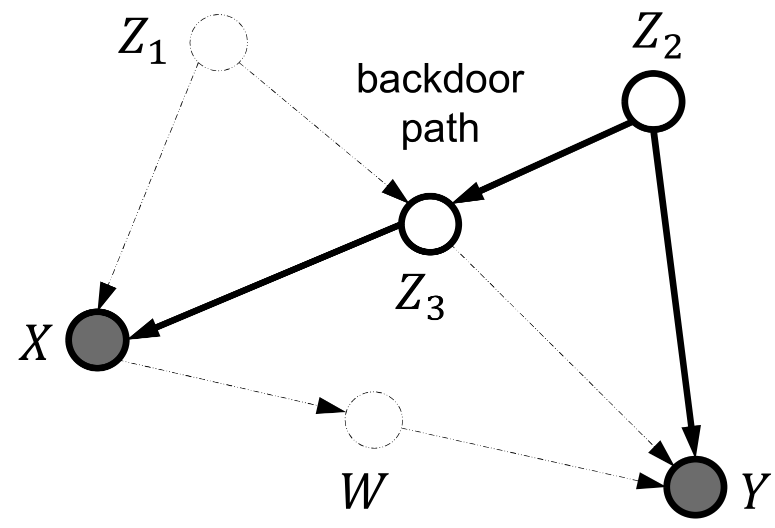

Figure 4.17

THE BACKDOOR CRITERION

Given an ordered pair of variables (X, Y) in a DAG $\mathcal{G}$, a set of variables $\mathbf{S}$ satisfies the backdoor criterion relative to (X, Y) if no node in $\mathbf{S}$ is a descendant of X, and $\mathbf{S}$ blocks every path between X and Y that contains an arrow into X.

The logic behind the **BACKDOOR CRITERION** is fairly straightforward.

In general, we would like to condition on a set of nodes $\mathbf{S}$ (**CONDITIONING SET**) such that we:

1. block all spurious paths between X and Y
2. leave all directed paths from X to Y unperturbed
3. create no new spurious paths

1. block all spurious paths between *X* and *Y*.

We want the conditioning set $\mathbf{S}$ to block any **BACKDOOR PATH** in which one end has an arrow into X, because such paths may make X and Y dependent, but are obviously not transmitting causal influences from X, and if we do not block them, they will confound the effect that X has on Y.

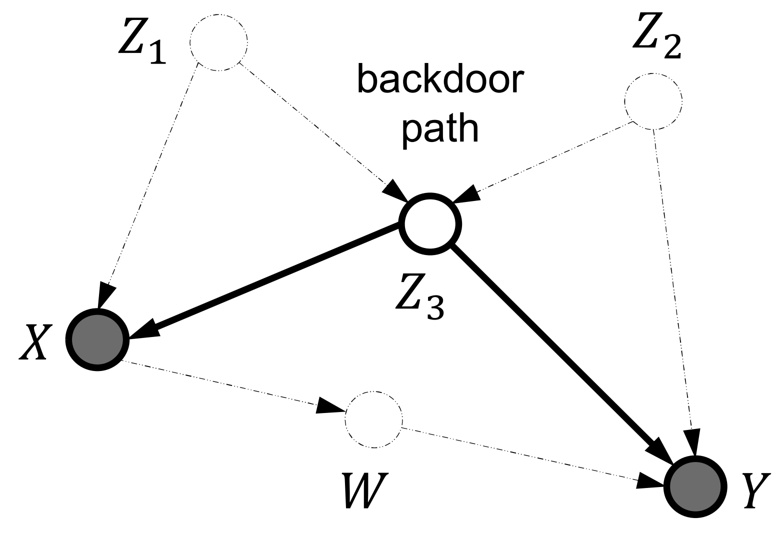

Figure 4.18

THE BACKDOOR CRITERION

Given an ordered pair of variables (X, Y) in a DAG $\mathcal{G}$, a set of variables $\mathbf{S}$ satisfies the backdoor criterion relative to (X, Y) if no node in $\mathbf{S}$ is a descendant of X, and $\mathbf{S}$ blocks every path between X and Y that contains an arrow into X.

The logic behind the **BACKDOOR CRITERION** is fairly straightforward.

In general, we would like to condition on a set of nodes $\mathbf{S}$ (**CONDITIONING SET**) such that we:

1. block all spurious paths between X and Y
2. leave all directed paths from X to Y unperturbed
3. create no new spurious paths

1. block all spurious paths between *X* and *Y*.

We want the conditioning set $\mathbf{S}$ to block any **BACKDOOR PATH** in which one end has an arrow into X, because such paths may make X and Y dependent, but are obviously not transmitting causal influences from X, and if we do not block them, they will confound the effect that X has on Y.

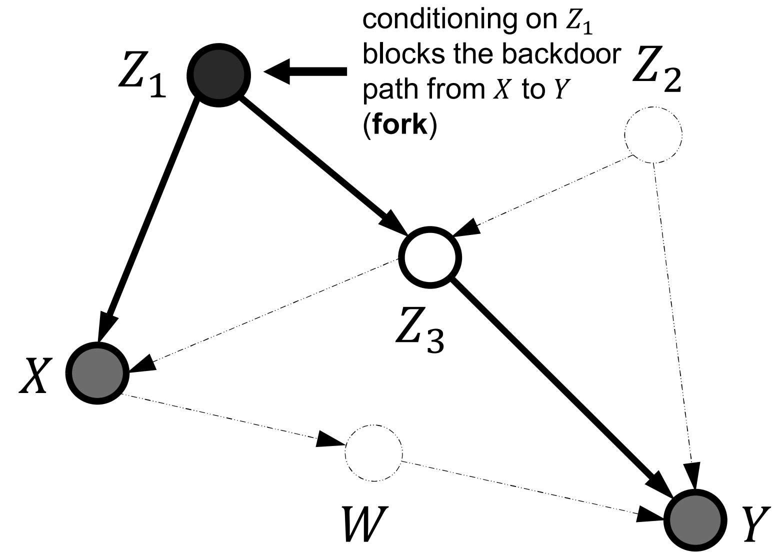

Figure 4.19

THE BACKDOOR CRITERION

Given an ordered pair of variables (X, Y) in a DAG $\mathcal{G}$, a set of variables $\mathbf{S}$ satisfies the backdoor criterion relative to (X, Y) if no node in $\mathbf{S}$ is a descendant of X, and $\mathbf{S}$ blocks every path between X and Y that contains an arrow into X.

The logic behind the **BACKDOOR CRITERION** is fairly straightforward.

In general, we would like to condition on a set of nodes $\mathbf{S}$ (**CONDITIONING SET**) such that we:

1. block all spurious paths between X and Y
2. leave all directed paths from X to Y unperturbed
3. create no new spurious paths

1. block all spurious paths between *X* and *Y*.

We want the conditioning set $\mathbf{S}$ to block any **BACKDOOR PATH** in which one end has an arrow into X, because such paths may make X and Y dependent, but are obviously not transmitting causal influences from X, and if we do not block them, they will confound the effect that X has on Y.

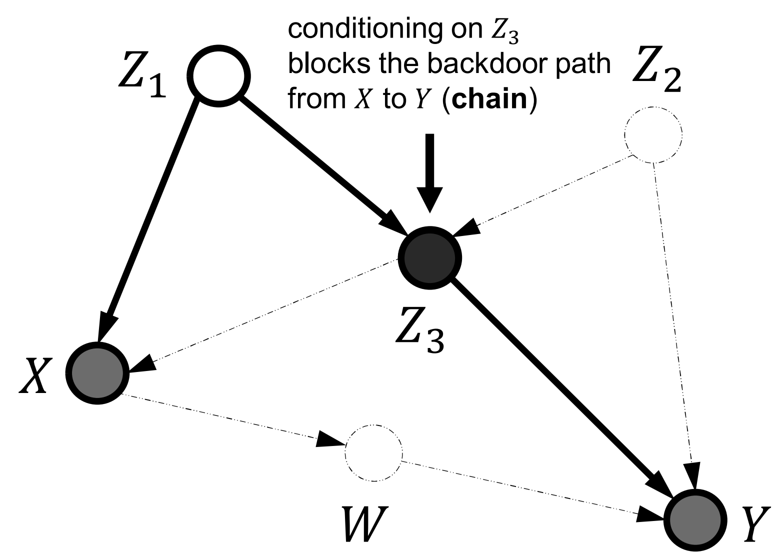

Figure 4.20

THE BACKDOOR CRITERION

Given an ordered pair of variables (X, Y) in a DAG $\mathcal{G}$, a set of variables $\mathbf{S}$ satisfies the backdoor criterion relative to (X, Y) if no node in $\mathbf{S}$ is a descendant of X, and $\mathbf{S}$ blocks every path between X and Y that contains an arrow into X.

The logic behind the **BACKDOOR CRITERION** is fairly straightforward.

In general, we would like to condition on a set of nodes $\mathbf{S}$ (**CONDITIONING SET**) such that we:

1. block all spurious paths between X and Y
2. leave all directed paths from X to Y unperturbed
3. create no new spurious paths

1. block all spurious paths between *X* and *Y*.

We want the conditioning set $\mathbf{S}$ to block any **BACKDOOR PATH** in which one end has an arrow into X, because such paths may make X and Y dependent, but are obviously not transmitting causal influences from X, and if we do not block them, they will confound the effect that X has on Y.

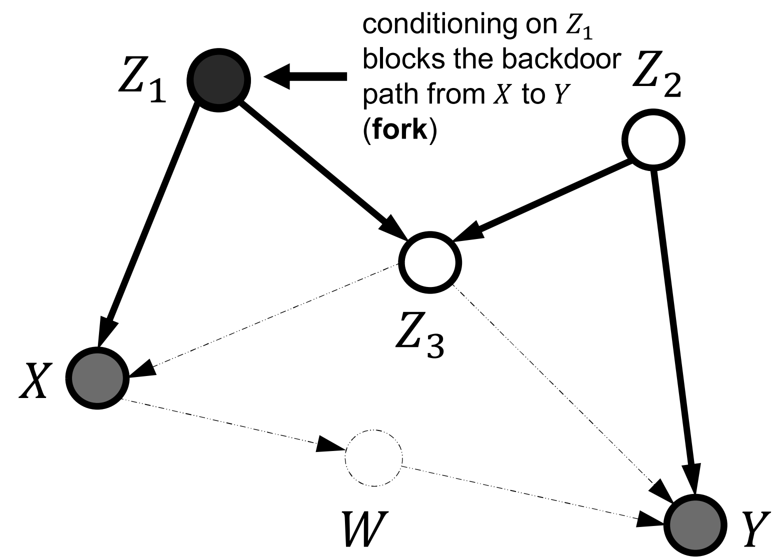

Figure 4.21

THE BACKDOOR CRITERION

Given an ordered pair of variables (X, Y) in a DAG $\mathcal{G}$, a set of variables $\mathbf{S}$ satisfies the backdoor criterion relative to (X, Y) if no node in $\mathbf{S}$ is a descendant of X, and $\mathbf{S}$ blocks every path between X and Y that contains an arrow into X.

The logic behind the **BACKDOOR CRITERION** is fairly straightforward.

In general, we would like to condition on a set of nodes $\mathbf{S}$ (**CONDITIONING SET**) such that we:

1. block all spurious paths between X and Y
2. leave all directed paths from X to Y unperturbed
3. create no new spurious paths

1. block all spurious paths between *X* and *Y*.

We want the conditioning set $\mathbf{S}$ to block any **BACKDOOR PATH** in which one end has an arrow into X, because such paths may make X and Y dependent, but are obviously not transmitting causal influences from X, and if we do not block them, they will confound the effect that X has on Y.

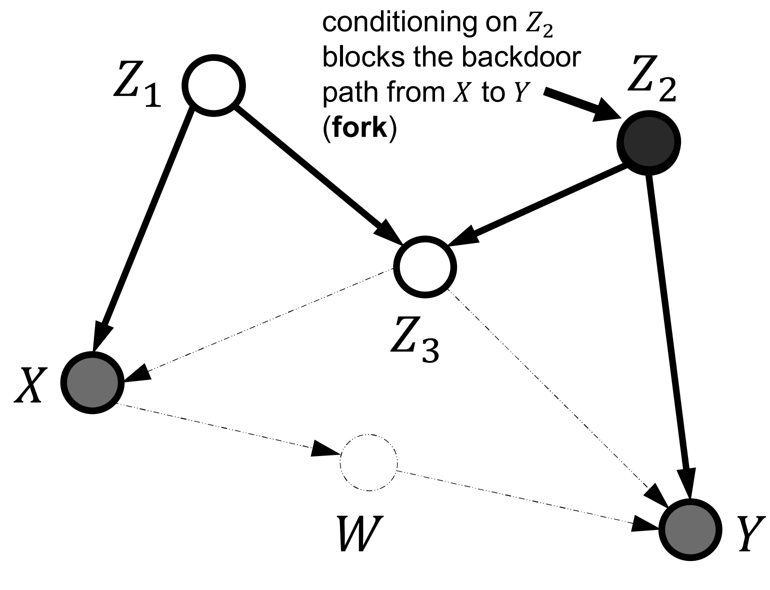

Figure 4.22

THE BACKDOOR CRITERION

Given an ordered pair of variables (X, Y) in a DAG $\mathcal{G}$, a set of variables $\mathbf{S}$ satisfies the backdoor criterion relative to (X, Y) if no node in $\mathbf{S}$ is a descendant of X, and $\mathbf{S}$ blocks every path between X and Y that contains an arrow into X.

The logic behind the **BACKDOOR CRITERION** is fairly straightforward.

In general, we would like to condition on a set of nodes $\mathbf{S}$ (**CONDITIONING SET**) such that we:

1. block all spurious paths between X and Y
2. leave all directed paths from X to Y unperturbed
3. create no new spurious paths

1. block all spurious paths between *X* and *Y*.

We want the conditioning set $\mathbf{S}$ to block any **BACKDOOR PATH** in which one end has an arrow into X, because such paths may make X and Y dependent, but are obviously not transmitting causal influences from X, and if we do not block them, they will confound the effect that X has on Y.

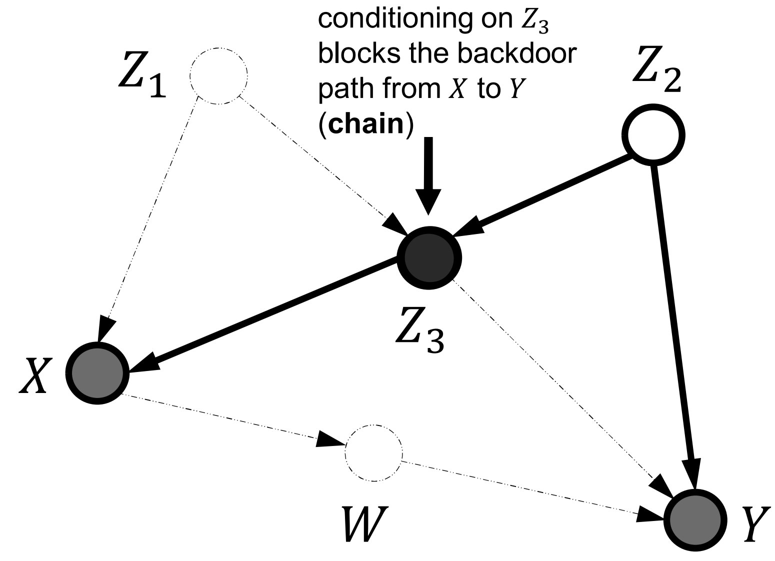

Figure 4.23

THE BACKDOOR CRITERION

Given an ordered pair of variables (X, Y) in a DAG $\mathcal{G}$, a set of variables $\mathbf{S}$ satisfies the backdoor criterion relative to (X, Y) if no node in $\mathbf{S}$ is a descendant of X, and $\mathbf{S}$ blocks every path between X and Y that contains an arrow into X.

The logic behind the **BACKDOOR CRITERION** is fairly straightforward.

In general, we would like to condition on a set of nodes $\mathbf{S}$ (**CONDITIONING SET**) such that we:

1. block all spurious paths between X and Y
2. leave all directed paths from X to Y unperturbed
3. create no new spurious paths

1. block all spurious paths between *X* and *Y*.

We want the conditioning set $\mathbf{S}$ to block any **BACKDOOR PATH** in which one end has an arrow into X, because such paths may make X and Y dependent, but are obviously not transmitting causal influences from X, and if we do not block them, they will confound the effect that X has on Y.

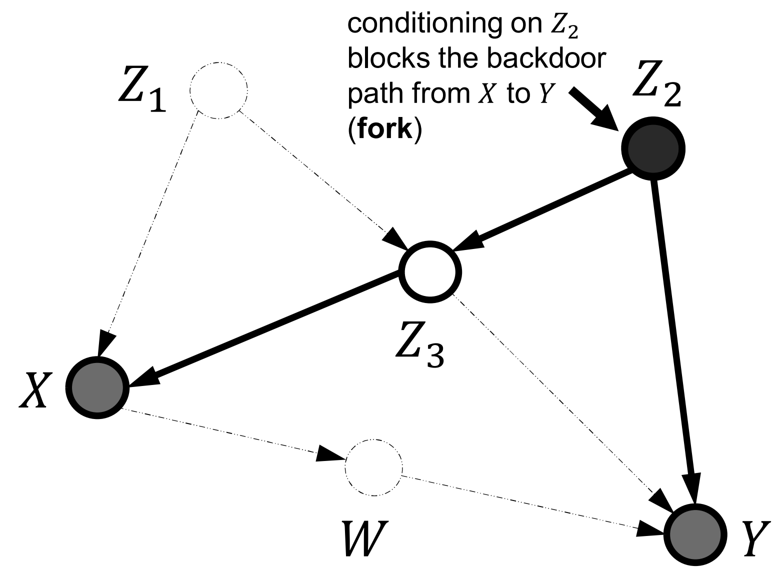

Figure 4.24

THE BACKDOOR CRITERION

Given an ordered pair of variables (X, Y) in a DAG $\mathcal{G}$, a set of variables $\mathbf{S}$ satisfies the backdoor criterion relative to (X, Y) if no node in $\mathbf{S}$ is a descendant of X, and $\mathbf{S}$ blocks every path between X and Y that contains an arrow into X.

The logic behind the **BACKDOOR CRITERION** is fairly straightforward.

In general, we would like to condition on a set of nodes $\mathbf{S}$ (**CONDITIONING SET**) such that we:

1. block all spurious paths between X and Y
2. leave all directed paths from X to Y unperturbed
3. create no new spurious paths

1. block all spurious paths between *X* and *Y*.

We want the conditioning set $\mathbf{S}$ to block any **BACKDOOR PATH** in which one end has an arrow into X, because such paths may make X and Y dependent, but are obviously not transmitting causal influences from X, and if we do not block them, they will confound the effect that X has on Y.

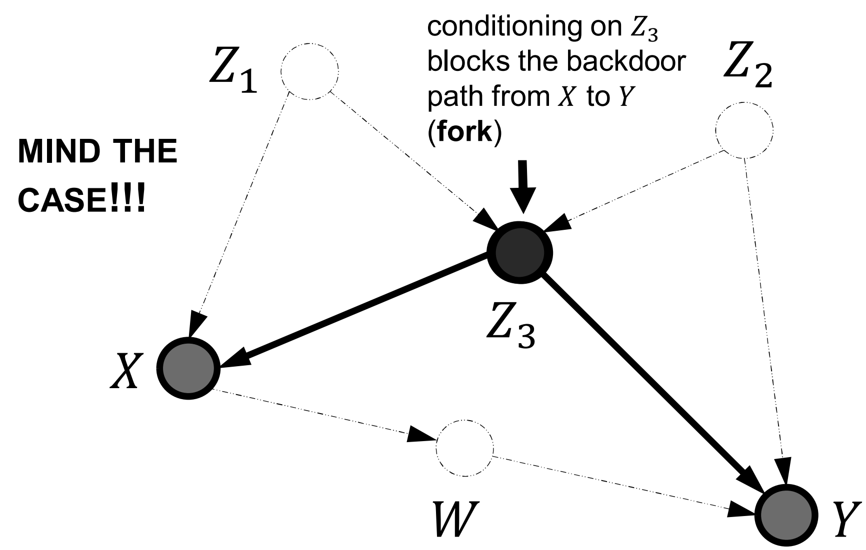

Figure 4.25

THE BACKDOOR CRITERION

Given an ordered pair of variables (X, Y) in a DAG $\mathcal{G}$, a set of variables $\mathbf{S}$ satisfies the backdoor criterion relative to (X, Y) if no node in $\mathbf{S}$ is a descendant of X, and $\mathbf{S}$ blocks every path between X and Y that contains an arrow into X.

The logic behind the **BACKDOOR CRITERION** is fairly straightforward.

In general, we would like to condition on a set of nodes $\mathbf{S}$ (**CONDITIONING SET**) such that we:

1. block all spurious paths between X and Y
2. leave all directed paths from X to Y unperturbed
3. create no new spurious paths

2. leave all directed paths from X to Y unperturbed

However, we don't want to condition on any nodes that are descendants of X.

Descendants of X would be affected by an intervention on X and might themselves affect Y; conditioning on them would block those pathways.

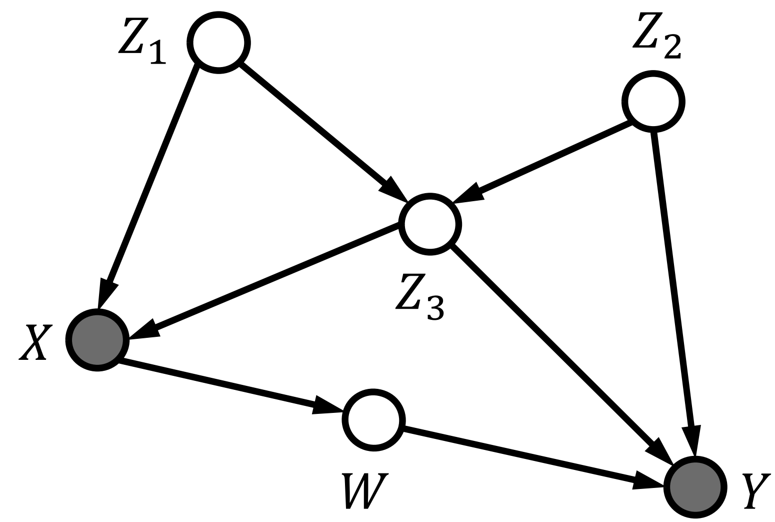

Figure 4.14

THE BACKDOOR CRITERION

Given an ordered pair of variables (X, Y) in a DAG $\mathcal{G}$, a set of variables $\mathbf{S}$ satisfies the backdoor criterion relative to (X, Y) if no node in $\mathbf{S}$ is a descendant of X, and $\mathbf{S}$ blocks every path between X and Y that contains an arrow into X.

The logic behind the **BACKDOOR CRITERION** is fairly straightforward.

In general, we would like to condition on a set of nodes $\mathbf{S}$ (**CONDITIONING SET**) such that we:

1. block all spurious paths between X and Y
2. leave all directed paths from X to Y unperturbed
3. create no new spurious paths

2. leave all directed paths from X to Y unperturbed

However, we don't want to condition on any nodes that are descendants of X.

Descendants of X would be affected by an intervention on X and might themselves affect Y; conditioning on them would block those pathways.

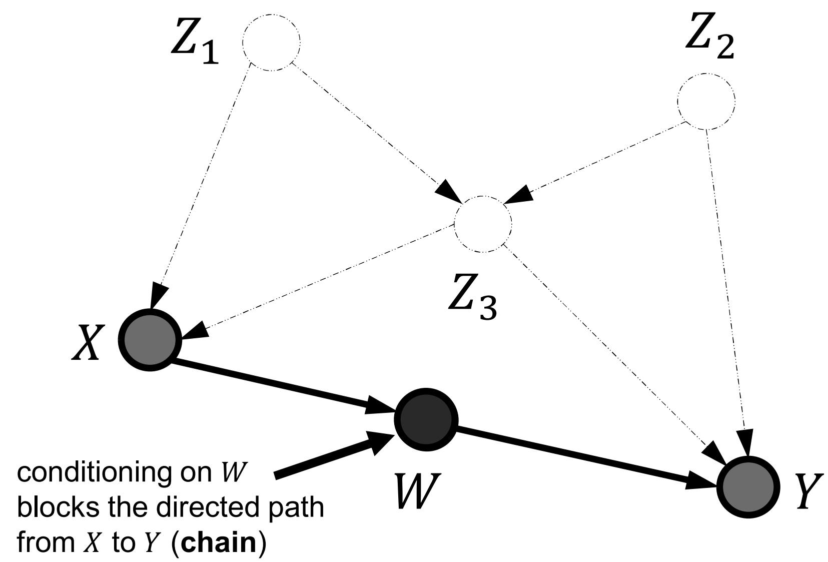

Figure 4.26

THE BACKDOOR CRITERION

Given an ordered pair of variables (X, Y) in a DAG $\mathcal{G}$, a set of variables $\mathbf{S}$ satisfies the backdoor criterion relative to (X, Y) if no node in $\mathbf{S}$ is a descendant of X, and $\mathbf{S}$ blocks every path between X and Y that contains an arrow into X.

The logic behind the **BACKDOOR CRITERION** is fairly straightforward.

In general, we would like to condition on a set of nodes $\mathbf{S}$ (**CONDITIONING SET**) such that we:

1. block all spurious paths between X and Y
2. leave all directed paths from X to Y unperturbed
3. create no new spurious paths

2. leave all directed paths from X to Y unperturbed

However, we don't want to condition on any nodes that are descendants of X.

Descendants of X would be affected by an intervention on X and might themselves affect Y; conditioning on them would block those pathways.

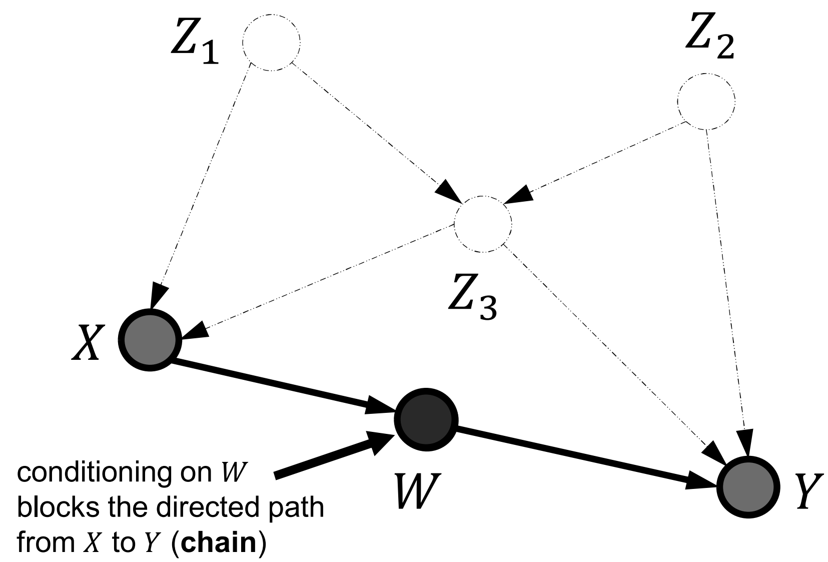

Figure 4.26

THE BACKDOOR CRITERION

Given an ordered pair of variables (X, Y) in a DAG $\mathcal{G}$, a set of variables $\mathbf{S}$ satisfies the backdoor criterion relative to (X, Y) if no node in $\mathbf{S}$ is a descendant of X, and $\mathbf{S}$ blocks every path between X and Y that contains an arrow into X.

The logic behind the **BACKDOOR CRITERION** is fairly straightforward.

In general, we would like to condition on a set of nodes $\mathbf{S}$ (**CONDITIONING SET**) such that we:

1. block all spurious paths between X and Y
2. leave all directed paths from X to Y unperturbed
3. create no new spurious paths

3. create no spurious paths

Finally, to comply with the third requirement, we should refrain from conditioning on any collider that would unblock a new path between X and X.

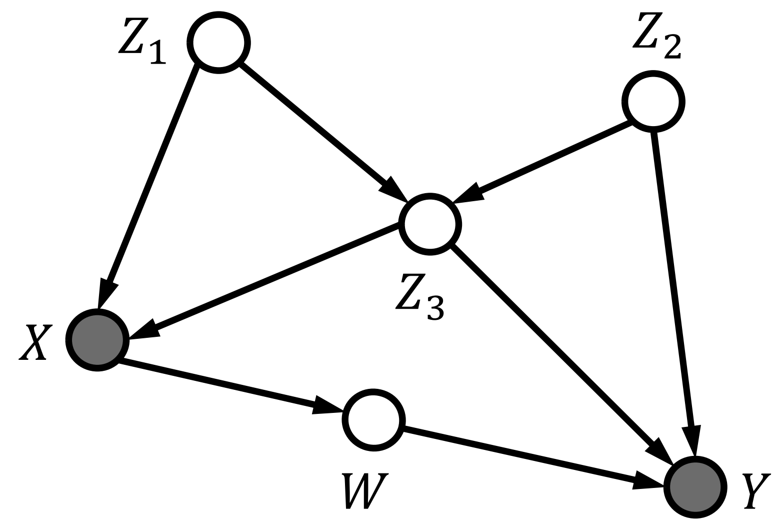

Figure 4.14

THE BACKDOOR CRITERION

Given an ordered pair of variables (X, Y) in a DAG $\mathcal{G}$, a set of variables $\mathbf{S}$ satisfies the backdoor criterion relative to (X, Y) if no node in $\mathbf{S}$ is a descendant of X, and $\mathbf{S}$ blocks every path between X and Y that contains an arrow into X.

The logic behind the **BACKDOOR CRITERION** is fairly straightforward.

In general, we would like to condition on a set of nodes $\mathbf{S}$ (**CONDITIONING SET**) such that we:

1. block all spurious paths between X and Y
2. leave all directed paths from X to Y unperturbed
3. create no new spurious paths

3. create no spurious paths

Finally, to comply with the third requirement, we should refrain from conditioning on any **COLLIDER** that would unblock a new path between X and X.

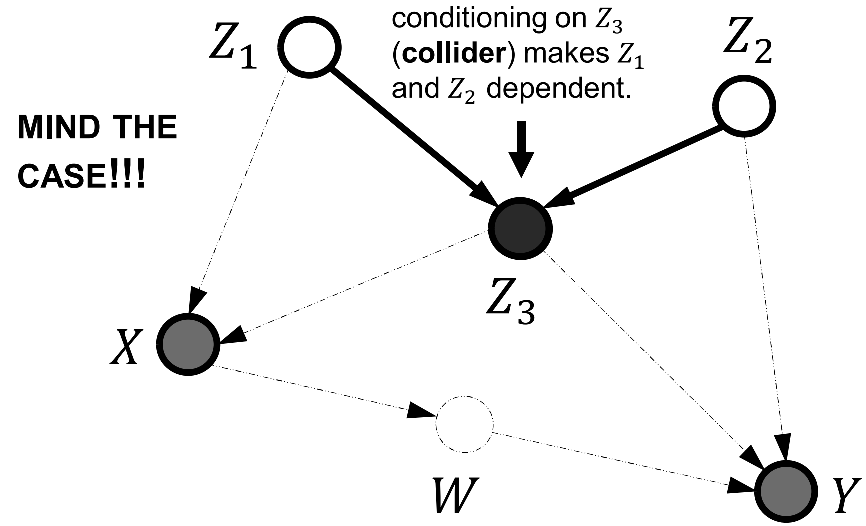

Figure 4.27

THE BACKDOOR CRITERION

Given an ordered pair of variables (X, Y) in a DAG $\mathcal{G}$, a set of variables $\mathbf{S}$ satisfies the backdoor criterion relative to (X, Y) if no node in $\mathbf{S}$ is a descendant of X, and $\mathbf{S}$ blocks every path between X and Y that contains an arrow into X.

The logic behind the **BACKDOOR CRITERION** is fairly straightforward.

In general, we would like to condition on a set of nodes $\mathbf{S}$ (**CONDITIONING SET**) such that we:

1. block all spurious paths between X and Y
2. leave all directed paths from X to Y unperturbed
3. create no new spurious paths

3. create no spurious paths

Finally, to comply with the third requirement, we should refrain from conditioning on any **COLLIDER** that would unblock a new path between X and X.

MIND THE CASE!!!

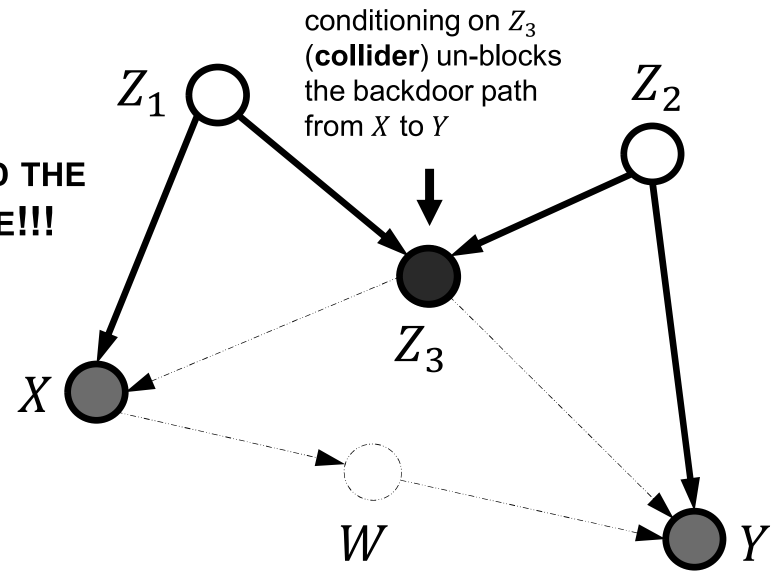

Figure 4.28

THE BACKDOOR CRITERION

Given an ordered pair of variables (X, Y) in a DAG $\mathcal{G}$, a set of variables $\mathbf{S}$ satisfies the backdoor criterion relative to (X, Y) if no node in $\mathbf{S}$ is a descendant of X, and $\mathbf{S}$ blocks every path between X and Y that contains an arrow into X.

The logic behind the **BACKDOOR CRITERION** is fairly straightforward.

In general, we would like to condition on a set of nodes $\mathbf{S}$ (**CONDITIONING SET**) such that we:

1. block all spurious paths between X and Y
2. leave all directed paths from X to Y unperturbed
3. create no new spurious paths

3. create no spurious paths

Finally, to comply with the third requirement, we should refrain from conditioning on any collider that would unblock a new path between X and X.

The requirement of excluding descendants of X, i.e. $de(X)$, also protects us from conditioning on children of intermediate nodes between X and Y (e.g., Q).

Figure 4.29

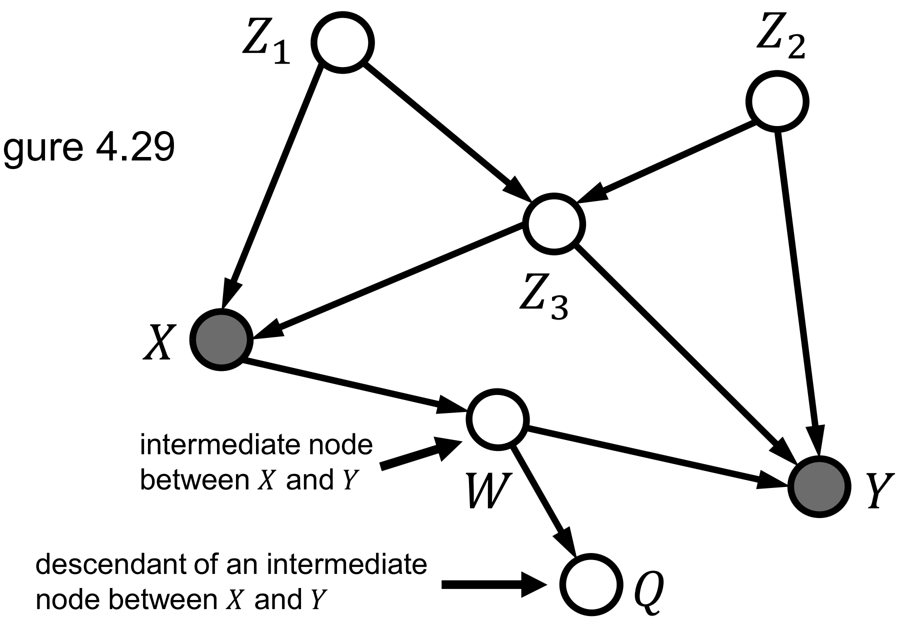

THE BACKDOOR CRITERION

Given an ordered pair of variables (X, Y) in a DAG $\mathcal{G}$, a set of variables $\mathbf{S}$ satisfies the backdoor criterion relative to (X, Y) if no node in $\mathbf{S}$ is a descendant of X, and $\mathbf{S}$ blocks every path between X and Y that contains an arrow into X.

The logic behind the **BACKDOOR CRITERION** is fairly straightforward.

In general, we would like to condition on a set of nodes $\mathbf{S}$ (**CONDITIONING SET**) such that we:

1. block all spurious paths between X and Y
2. leave all directed paths from X to Y unperturbed
3. create no new spurious paths

3. create no spurious paths

Finally, to comply with the third requirement, we should refrain from conditioning on any collider that would unblock a new path between X and X.

The requirement of excluding descendants of X, i.e. $de(X)$, also protects us from conditioning on children of intermediate nodes between X and Y (e.g., Q).

Figure 4.30

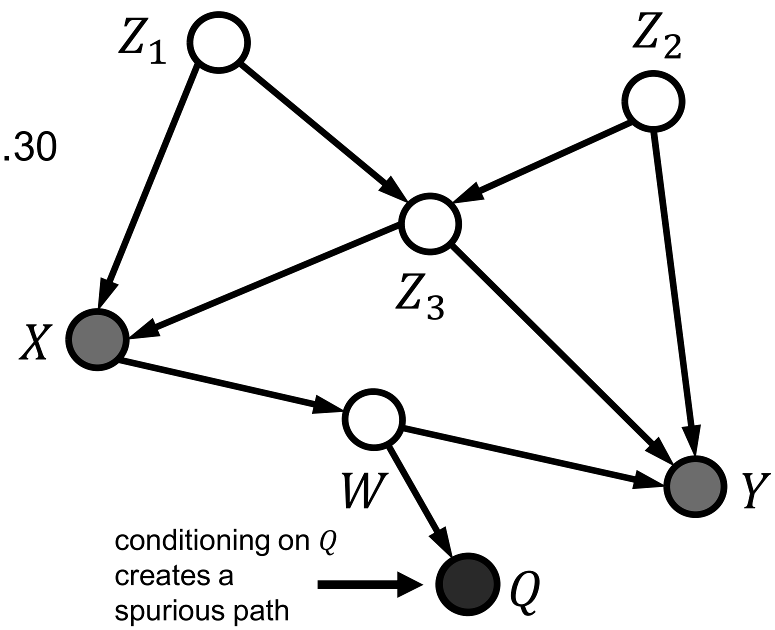

THE BACKDOOR CRITERION

Given an ordered pair of variables (X, Y) in a DAG $\mathcal{G}$, a set of variables $\mathbf{S}$ satisfies the backdoor criterion relative to (X, Y) if no node in $\mathbf{S}$ is a descendant of X, and $\mathbf{S}$ blocks every path between X and Y that contains an arrow into X.

The logic behind the **BACKDOOR CRITERION** is fairly straightforward.

In general, we would like to condition on a set of nodes $\mathbf{S}$ (**CONDITIONING SET**) such that we:

1. block all spurious paths between X and Y
2. leave all directed paths from X to Y unperturbed
3. create no new spurious paths

3. create no spurious paths

Finally, to comply with the third requirement, we should refrain from conditioning on any collider that would unblock a new path between X and X.

The requirement of excluding descendants of X, i.e. $de(X)$, also protects us from conditioning on children of intermediate nodes between X and Y (e.g., Q).

Figure 4.31

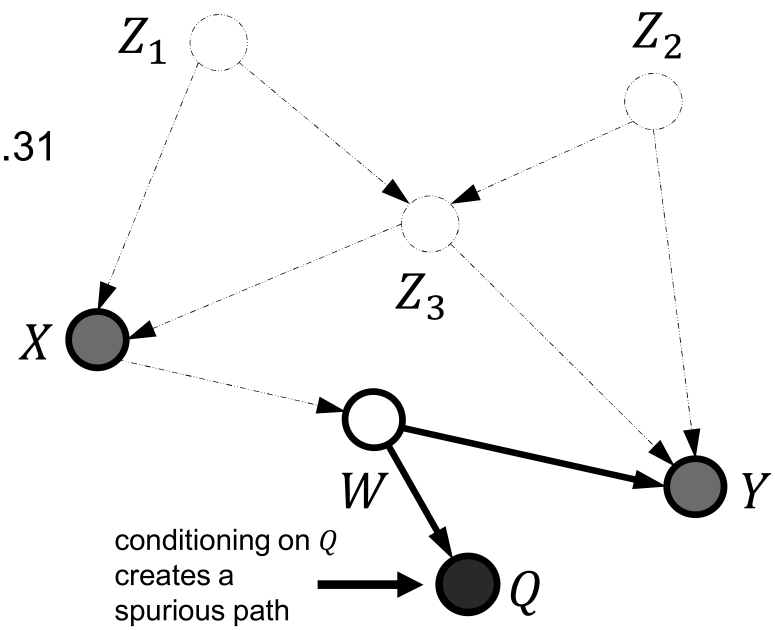

THE BACKDOOR CRITERION

Given an ordered pair of variables (X, Y) in a DAG $\mathcal{G}$, a set of variables $\mathbf{S}$ satisfies the backdoor criterion relative to (X, Y) if no node in $\mathbf{S}$ is a descendant of X, and $\mathbf{S}$ blocks every path between X and Y that contains an arrow into X.

THE BACKDOOR ADJUSTMENT FORMULA

Given the modularity assumption, that, $\mathbf{S}$ satisfies the backdoor criterion, and positivity, we can identify the causal effect of X on Y as follows:

$$P\big(Y = y|do(X = x)\big) = \sum_{\boldsymbol{s}} P(Y = y|X = x, \mathbf{S} = \boldsymbol{s})\ P(\mathbf{S} = \boldsymbol{s})$$

$$P\big(Y = y|do(X = x)\big) = \sum_{\boldsymbol{s}} P(Y = y|do(X = x), \mathbf{S} = \boldsymbol{s})\ P\big(\mathbf{S} = \boldsymbol{s}|do(X = x)\big)$$

(conditioning on $\mathbf{S}$ and then marginalizing $\mathbf{S}$ out)

$$= \sum_{\boldsymbol{s}} P(Y = y|X = x, \mathbf{S} = \boldsymbol{s})\ P\big(\mathbf{S} = \boldsymbol{s}|do(X = x)\big)$$

($\mathbf{S}$ satisfies the backdoor criterion)

$$= \sum_{\boldsymbol{s}} P(Y = y|X = x, \mathbf{S} = \boldsymbol{s})\ P(\mathbf{S} = \boldsymbol{s})$$

(by modularity)

A set $\mathbf{S}$ which satisfies the backdoor criterion is said to be a **SUFFICIENT ADJUSTMENT SET**.

THE BACKDOOR CRITERION

Given an ordered pair of variables (X, Y) in a DAG $\mathcal{G}$, a set of variables $\mathbf{S}$ satisfies the backdoor criterion relative to (X, Y) if no node in $\mathbf{S}$ is a descendant of X, and $\mathbf{S}$ blocks every path between X and Y that contains an arrow into X.

THE BACKDOOR ADJUSTMENT FORMULA

Given the modularity assumption, that, $\mathbf{S}$ satisfies the backdoor criterion, and positivity, we can identify the causal effect of X on Y as follows:

$$P\big(Y = y|do(X = x)\big) = \sum_{\boldsymbol{s}} P(Y = y|X = x, \mathbf{S} = \boldsymbol{s})\, P(\mathbf{S} = \boldsymbol{s})$$

We can use the backdoor adjustment formula if, $\mathbf{S}$ *d*-separates X from Y in the **AUGMENTED GRAPH** (obtained by removing all outgoing edges from X).

In previous lectures we mentioned that we would be able to isolate the causal association if X is *d*-separated from Y in the **AUGMENTED GRAPH.**

"Isolation of the causal association" is identification.

We can also isolate the causal association if X is *d*-separated from Y in the **AUGMENTED GRAPH**, conditional on $\mathbf{S}$.

This is what the first part of the backdoor criterion is about and what we've codified in the backdoor adjustment.

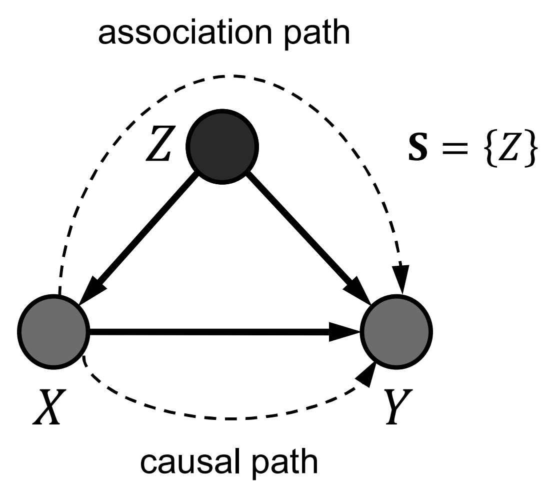

Figure 4.32